

Response to the referee 1 for “The ESCAPE project: Energy-efficient Scalable Algorithms for Weather Prediction at Exascale”

(grey background: text of the reviewer comment, white background: our response)

Overall impression

The manuscript under consideration was submitted to GMD as a "development and technical paper". The paper topically spans mathematical formulation, numerical integration techniques, parallelization strategies, language-specific aspects of implementation, hardware-specific optimizations, hardware construction, and operational considerations in the context of both global-circulation and limited-area models.

As the main conclusion from the hereby review, I propose to significantly shorten the article (currently 50-pages and 30 figures) and change its type to a "Review and perspective paper" to match the stated intention of the authors to create “the flagship publication for the EU project ESCAPE ... introduce the concept of weather & climate dwarfs and discuss first results in terms of optimization and performance portability”. This path seems to me as much more reasonable than working towards matching the requirements for a “development and technical paper” as defined by GMD guidelines.

The authors aim at: (i) providing a technical report touching upon current-hardwarespecific performance measures, (ii) structuring the work as a research paper, and (iii) presenting a project-promoting overview article. These aims are incompatible in my opinion, and trying to achieve all of them at once results in unclear target audience and an apparent lack of a storyline, despite high potential for strong conclusions to be based on the presented results.

Notwithstanding, I do see a point in publishing such a “perspective” paper with the aim of promoting the project results and giving due credit to participating parties. I expect such a shorter “perspective” paper to achieve a higher impact and I encourage the editorial team to offer this option to the authors.

We would like to thank the reviewer for this advice. Our first submission attempted to describe all the work that was done in the ESCAPE project. We agree that this was too much material and distracted from the main message of the paper. We have significantly shortened the paper and focus now on properly motivating the dwarf concept and to demonstrate this concept by describing the work for one of the dwarfs in detail.

Code availability

The “Code availability” section on page 42 is derisory. The standard that GMD is fostering among the community is to enable readers and reviewers to reproduce results presented in GMD papers. Here, the reader is only given a link to project website where one may not even find a properly defined software license – just a statement that it "permits free of charge use for educational and/or non-commercial research". The final sentence of the referenced website reads: “If you wish to access any of the implementations, please contact us via the contact form and we will provide further information on the process of obtaining a license”. This stands in clear opposition to the anonymous public access recommendations of GMD. Basing on an educated guess (the most one can anonymously base on given the above), I consider the results presented in the paper as not independently reproducible for reasons including software and hardware availability, as well as lack of availability of the details of the test cases.

As outlined above, a solution would be to move much of the technical details to another publication (a technical report issued by one of the participating institutions – several of those are already cited) and present a “perspective” paper for which GMD does offer an exception in

terms of reproducibility level (see https://www.geoscientific-model-development.net/about/manuscript_types.html).

The rules in terms of code availability and reproducibility for a development and technical paper are according to the link given by the reviewer the same as for a model description paper. The website of GMD states that "When copyright or licensing restrictions prevent the public release of model code, or in the cases where there is some other good reason for not allowing public access to the code, topical editors must still be given access to the model code. Access must also be granted to the reviewers whilst preserving their anonymity, if this is legally possible." We are happy to offer the editor and the reviewers a license to access the code and test cases used in our work. There is no need to make an exception for this paper in terms of reproducibility.

Nevertheless, the code availability section should contain comprehensive information, and clearly inform about the code availability, not license availability. Please point to repositories, state precisely the licenses or clearly indicate if the code is not publicly available or its reuse is constrained. In case of lack of public availability, GMD requires to state the reasons for it. Please include information for all the software that was essential in obtaining the presented results, including the participating weather prediction models: IFS, ALARO, COSMO-EULAG as well as the described tools such as GRASS, CLAW and GridTools (the <https://github.com/eth-cscs/gridtools> repository linked from the GridTools website does not exist as of time of writing this review).

We have changed the license statement and added information about all software used in the ESCAPE project.

Code availability for hardware-specific tools such as those essential in GPU code development should also be included in the section (see doi:10.1002/2016WR020190 for a recent discussion in the context of hydrological modelling), and in my opinion should also be included in the discussion if a proper "perspective" is to be given. The paper gives an overview of several paths forward in NWP systems development and aims at discussing longer-term strategies. Such discussion calls for mentioning which optimisation strategies are prone to the vendor lock-in threat.

We have added information about the compilers and tools that we used and their versions. As described in the paper we explore vendor specific strategies but we also aim at avoiding vendor lock-in through the use of domain specific languages.

The dwarf nomenclature and technicalities

The authors highlight throughout the paper the concept of separation of concerns in software engineering using the notion of a "dwarf" which the paper introduces in the context of weather and climate models. The reason to introduce a new term is not given. What does the new concept replace (monoliths)? The adopted term is seemingly wrongly attributed (Colella as opposed to Asanović et al.?) and, in my understanding, used in a misleading way. The reason why the 7 dwarfs of Colella, and later the 13 Berkeley dwarfs, the 7 dwarfs of Symbolic Computation, the 13 Parallel Dwarfs (and likely others) were introduced is that the concept they generalize does not easily fit into existing encapsulation nomenclature of: components, frameworks, layers, subsystems, libraries, kernels, modules, services, drivers, plug-ins, controllers, etc. Why dynamical core layers, physics modules and numerical libraries are to be renamed? In principle, why not – let us embrace the introduced notion of Weather & Climate Dwarfs, but please do clarify in the paper the reasons to introduce the new nomenclature and clearly differentiate it from existing solutions.

We referred to Asanović et al. because we could not find a good reference for Colella's presentation. We have revised the motivation section for the dwarf concept and we have added new references.

Moreover, among the community, dwarfs are being defined as computation/communication idioms (e.g., OpenDwarfs project) or (design) patterns of high performance computing (e.g., the mentioned work of Asanović) – is it not misleading to refer to the “components of an Earth system model” borrowing the term that was introduced to define idioms/patterns?

Our goal is to identify patterns in terms of computation and communication that are characteristic for weather and climate models. Each of our dwarfs possesses very characteristic patterns of high performance computing. We have added a column to Table 1 describing the characteristics of each dwarf in terms of communication and computation.

In the conclusions section, the authors note that “... the entire NWP community has a strong interest in pursuing the creation of dwarfs”. Such community-wide initiatives are, in my opinion, only feasible if backed by interface standardization and “liberation” of the software in question. Please do elaborate on how (technically and legally) the reusability of the envisaged community “dwarfs” is to be assured? Furthermore, the dwarfs’ interfaces will determine if the sought after capability of overlapping computations will be feasible.

We have revised the quoted part of the conclusions. We agree from a technical point on clean interfaces, but scientifically disagree that standard interfaces are useful (i.e. plug and play concepts of specific parametrisations or physics and dynamics). In our view standardisation of interfaces in operational NWP software, despite being a very attractive technical concept, it is often detrimental to forecast quality and computational efficiency. It is not the aim of ESCAPE to develop standardised interfaces, even if we explore overlapping concepts. In fact bespoke interfacing may become more important with increasing resolution.

Other remarks

- Optimization is the leitmotif of the paper, and is exemplified with detailed results obtained with current hardware and described in zealous detail (a somehow anecdotal example: replace division by constant with multiplication by the reciprocal calculated in advance). Yet at the same time, the paper aims at providing “strategy to evolve weather and climate prediction models to next-generation computing technologies”, exascale is mentioned in the title. Please extend the conclusions section by indicating which of the employed optimizations have the chance to offer “software sustainability” over the years (also in the context of software maintainability tradeoffs), and where, e.g., indirection would be the key to performance (DSLs)?

We have added new results about running the spectral transform dwarf on the supercomputer Summit (currently the fastest supercomputer in the world) and we have added a discussion about the sustainability of the chosen techniques.

- There is an abundance of arcane nomenclature in the paper used without introduction and spanning a broad range of domains, e.g.: “loop fusion optimization”, “generic reduced Gaussian grid”, “parallelogram”, “a perturbative method which determines the (constrained) optical phase”, “Spherical Harmonics TCo639 test case”, “baroclinic instability benchmark”, “weighted line Jacobi method”, “V-cycle configuration”, “major generation and destruction processes, including cloud formation by detrainment from cumulus convection”;

We would like to thank the reviewer for identifying these issues. Many of these statements have been removed by shortening the paper. We have revised the remaining statements.

- There is an imbalance in the level of detail of different sections of the paper, e.g. the well-established MPDATA is introduced with an outline touching upon numerical analysis, while the “created” “global shallow-water model named GRASS” is presented with just a

reference to a pair of submitted papers; there is a two-pagelength introduction to the application of spectral transform methods in NWP, while less than that is devoted to the entire discussion of physics dwarfs; authors do admit that this was intentional (p. 5 last paragraph) for some of the work has not been published elsewhere – this however rather supports publishing it elsewhere than sneaking into an overview paper.

It was never our goal to cover every dwarf in detail. Our goal was to demonstrate our optimisation workflow for a small selection and briefly describe the others. To make this clearer we have now removed the dedicated subsections for the dwarfs which we did not cover in detail. Instead we refer to the corresponding publications in the dwarf table.

- Some of the tools mentioned target Fortran development (e.g., CLAW DSL) while other cater to a wider set of technologies (e.g., Atlas), this is not mentioned explicitly and the reader is left without a clear statement if the proposed directions of development deviate or not from the Fortran ecosystem;

All of the work presented is compatible with Fortran. Even the GPU optimisation is mostly done with OpenACC in Fortran. Gridtools requires currently C++ but we plan to provide Fortran support in future DSLs. We have added a statement about this in the paper.

- The conclusions section contains statements of overly contrasting time horizons: on the one hand, the authors mention “adding a large number of zero operations” what is explained in the text to be caused simply by lack of support for a particular feature in the current version of a third-party library; on the other hand, prerequisites and challenges for subkilometer global simulations are mentioned. Please reconsider what are the main project conclusions worth to be listed in the concluding section and abstract.

We have revised the conclusions section.

- That the great majority of referenced works is [co]authored by the manuscript authors amplifies the feeling of some of the methodology, design or vendor choices being given without a proper context on the alternatives:
 - How representative is the chosen set of models (IFS, ALARO, ALADIN and COSMO-EULAG) among the “competition” and how the considered speedup techniques compare with what has been explored recently (see, e.g., doi:10.1175/BAMS-D-15-00278.1 and references therein)?
 - How the proprietary software and hardware solutions like cuBLAS/cuFFT and NVLink/NVSwitch compare to those provided by other vendors?
 - Overlapping CPU-GPU computation strategy for dynamics/physics has been recently discussed in GMD in context of cloud-resolving simulations (doi:10.5194/gmd-2018-281, e.g. fig. Fig. 1), could the discussion here be supported with references to existing solutions from other domains?
 - mentions of GPU-resident weather forecasting call for citing other recent works (e.g., doi:10.1175/BAMS-D-14-00114.1)
 - What are the alternatives for the used radiation and cloud-physics schemes, are the chosen ones representative of what the community envisages for the mentioned global subkilometer-scale future simulations?
 - Are the CLAW and GridTools DSLs the sole solution available in this context?

We have added some references to the papers mentioned by the reviewer. A comparison with solutions by vendors not involved in the project is outside the scope of this paper. We

are not able and it is not our goal to provide an exhaustive discussion of all available strategies. The goal of the paper is to present the dwarf concept to the weather and climate community and to demonstrate through a few examples how it can be used to enable close collaboration between NWP centres and hardware vendors.

- The word “code” is used in a somehow casual way, e.g. “redesign of the algorithms and codes”, “our work on optimising codes”, “code used for data assimilation”, “models from which the dwarf code originated”; let me suggest to consider employing more cross-domain notions of implementation, software, etc; similar nomenclature issue: restructure vs. refactor;

We followed the reviewers advice, checked every instance of the word "code" and replaced the nomenclature where suitable.

- Please remove any mentions of internal labels within the code – this information is unneeded for a research paper audience: “halo_exchange subroutine”, “compute_fluxzdiv”, “this%geom%node2edge_sign”.

We have removed the internal labels.

- Please limit the use of acronyms/short-forms, and remove those clearly unneeded: PSNC in Fig. 3, EBTI on page 24, GP_dynamics/SP_transforms/SI_solver/RAD in caption of Fig. 5, semi-Lagrange in Fig. 6; Some references are listed with DOI number, some without - please be consistent; FORTRAN/Fortran, TRAP2/Trap2 spelling - please be consistent; Among the affiliations listed, some are given with detailed street addresses, some without - please be consistent.

We have removed the mentioned acronyms and made the text and affiliations consistent. The DOI numbers are given whenever they are available.

- The title of the paper reads “The ESCAPE project: Energy-efficient Scalable Algorithms for Weather Prediction at Exascale”. Exascale is not discussed or defined and barely mentioned only in the conclusions, while the phrase “weather & climate” is used throughout the paper.

We have added recent results on Summit and a discussion which is more targeted at exascale. The title of the paper is the name of the project.

- Statements such as “ECMWF is world leading in terms of track forecast”, “extreme computational capabilities typically required in operational forecast production”, “[IFS code] has been continuously optimized over multiple decades”, “Feedback from the European and international community at our dissemination workshops and at international conferences has shown that this work was well received” are, in my opinion, good candidates for removal when shortening the paper – please avoid promotional language and statements which are not falsifiable; another candidates for removal are numerous vague statements: “most speedup seems to be due to avoiding some of the temporary arrays”, “some more fundamental changes which are more difficult to apply”, “whole cycle might employ some form of smoother/solver”, “has to be wisely chosen according to the cluster hardware”, “we do not know if there will be a clear winner”, “The first results of this effort look promising”.

Many of these statements were removed in the process of shortening the paper as suggested by all reviewers. We have revised the remaining statements.

Response to the referee 2 for “The ESCAPE project: Energy-efficient Scalable Algorithms for Weather Prediction at Exascale”

(grey background: text of the reviewer comment, white background: our response)

The paper presents a review of the work done in the Escape project. This reviewer is familiar with some of this work and reporting the work done as part of the project is certainly of interest to the community. Unfortunately, the paper is not well written, it is full of mistakes, informal language and confusing or unclear explanations. I have read and documented changes as far as page 17, but this has taken a long time as the paper has not been properly proof read before submission. Referee #1 calls for a substantial revision, and a possible change of paper type, therefore there doesn't seem much point in fully detailing necessary changes beyond this point. A shorter, more focused article as a review and perspective paper would improve the readability and is probably more appropriate for the content.

We would like to thank reviewer 2 for his advice. We have significantly shortened the paper. The main purpose of the paper is to present the dwarf concept to the weather & climate community and to demonstrate it with a detailed example.

I include the detailed points below, which need to be addressed.

Page 2 line 7 Weather prediction (models?)

We have changed the sentence.

Line 20, sentence reads as if heavy precipitation patterns could lead to tropical cyclones, modify

We have revised the introduction.

Line 21 being satisfying “being”

We have changed this sentence.

Intro 1st paragraph is rather clumsy, there are plenty of reasons improved forecasts in general would have economic and societal benefit besides heavy precipitation. The need to improve resolution is given as the main motivation for improved forecasts but then Climate is thrown into the following sentence. Improved resolution versus complexity for improved climate is a matter of debate. There is no mention of current resolutions for the reader to compare 1km. What does global resolution range mean? The last sentence is also confusingly written. This paragraph needs re-writing with proper thought on what is the motivation for improving resolution of weather and climate simulations. There are plenty of justifications.

We have revised the introduction.

Line 27 “guarantee the continued efficiency” is probably a bit strong. “Enable efficient implementations of “ or similar is probably more realistic.

We have made the suggested replacement.

Line 33 “The authors there” is rather informal language.

We have changed this sentence.

Page 3. Line 7 Citation needed to back authors claim that ECMWF is world leading in terms of track forecast. It would be quite odd for a higher resolution forecast to “compromise the accuracy” of a lower resolution, especially as this is the motivation for escape.

We have revised this part of the introduction.

The paragraph on supercomputers is also rather clumsily written and there should be a reference. An obvious one would be “Crossing the Chasm : how to develop weather and climate models for next generation computers?”. Some of the authors of this paper are also authors of the chasm paper. Other references would also be beneficial.

We have revised the introduction and we have added the suggested reference.

Section 2, line 18, The use of “at once” suggests either “immediately” or “at the same time” but “too big” implies neither. This sentence needs re-writing.

We have changed this sentence.

Page 4 The text in figure 1 is far too small to read.

The small text in figure 1 contained only license terms of material that was used to create the figure. We tried to minimise the distraction from the message of the figure by making the text very small. We have now made the text larger to enable everyone to read the text.

Figure 2a “halo’s” has no apostrophe.

We removed this figure from the paper in order to shorten the paper.

Page 5 figure 3 What does the vertical line denote? Line 5 what does “released dwarf” mean?

We have added an explanation of the vertical line to the caption and we have revised the text.

Page 7 line 5 “use 3D var like the “ is too informal language, use for example or e.g.

We have removed this text in order to shorten the paper.

Page 7 line 23 “need” -> “needs”

We have removed this text in order to shorten the paper.

Page 8 Figure 8. The figures are not very well explained. A careful reading of a technical appendix is required and even then it is not clear how they illustrate the point that overlapping

the comms of data transfer is necessary. The main constraint for NWP is time to completion, it is not obvious that the assumptions (necessarily) made to express the scaling in terms of energy are sufficiently universal to infer the point the authors wish to make. Whilst using less energy reduces the cost, the run-time constraint is the primary motivation for NWP. This section needs to be re-worked to describe the performance model and explain how any figures used enable the authors draw the conclusions that they do.

We have removed the cost model in the process of shortening the paper. Just for clarification: all the curves in the plots of the cost model were created under the constraint that the forecast needs to finish within the required runtime. The cost model attempted to answer the question: under the given runtime constraint how much code needs to run on an accelerator to make the use of accelerators financially beneficial. As mentioned in the first version of the paper this cost model requires many assumptions and should only be used as a motivation to port as much of the model as possible to the accelerator. It may not be used as a basis for management decisions.

Page 9 section 3 line 23 Last sentence. What does this mean?

The inverse transform uses first an inverse Legendre transform to compute Fourier coefficients and applies then an FFT to obtain grid point values. The direct transform uses the opposite order: first an FFT to compute Fourier coefficients and then a direct Legendre transform to compute the spectral coefficients. We agree that the term "opposite direction" to describe this difference was misleading. We have revised this part of the text.

Page 10, refers to figure 5 which is on page 11. How are the profiles produced? What machine are they produced from, especially the node count etc would change the profile? The "Anticipated" future are these profiled from data or some performance model? The 1.25km run may well have a different profile on a different machine?

All of these profiles have been measured with the actual model on the Cray XC40 supercomputer of ECMWF. The number of nodes has been chosen such that the model run would be suitable for operational requirements. We agree that different machines will produce different profiles. These profiles are highly relevant for ECMWF because they represent the situation on the machine that is currently used for the operational forecast.

Page 11 refers to figure 7 which appears on page 13, the text "like on the left" is rather informal and inappropriate for a scientific paper.

We have changed the figure accordingly.

Page 12 refers to figure 8 which appears on page 14. In the figure, there is no open diamond referred to in the legend and the caption, the open rectangle referred to in the caption doesn't appear but a dash or line does which isn't referred to in the caption.

The open diamond of the legend is half covered by the filled circle in the plot. The open rectangle of the legend was an oversight and should have been the dash. We have made the filled circle open and changed the colour to make the overlapping points easier to recognise.

Pages 12 and 13 there is a complicated discussion of code changes. This would be illuminated by some code fragments as examples.

We have added code examples in section 3.3.

Figure 10. Page 16 The data points are connected with a line (something spreadsheet application does readily), however, the horizontal axis is Number of GPUs, which is a discrete variable, so a line graph is wrong. Whilst it is not unusual to see such a plot, the authors are not predicting the speed up on 16.5 GPUs so why the line? The plot should be re-drawn appropriately.

We agree with the reviewer that the connecting line between the data points is not supposed to provide a scientific message. We believe that the connecting line makes the plot more readable. We have reduced the thickness of the line by 50% to make it less prominent and we added a statement to the caption describing that the data points were connected with lines purely for the purpose of improving readability.

Page 17 lines 6-15 the paragraph discusses the pack-unpack operation. The sentence sender and receiver share their memory layout as the way they may differ is confusing. How can they share a memory layout if they are different? What was the change that made the performance improvement? Were the pack and unpack scanning memory unnecessarily?

We have replaced "share" with "exchange" in the sense that sender and receiver exchange their own layout among each other. The performance improvement came from reordering the loops for both pack and unpack following the memory layout of the scattered buffer. This optimisation decreased the number of tests (i.e. copy or not copy) and avoids scanning memory multiple times. Scanning the memory multiple times was unnecessary.

The final paragraph on page 17 is a discussion of implementing some of the GPU optimisations on CPUs. Again some code fragments here would be helpful. If the GPU optimisations are too intensive, how can they be used elsewhere for other architectures? The stated goal of ESCAPE is to re-assemble the models from optimised dwarfs. How is this managed? If there are conflicting optimisations how are they resolved? Is single source code possible?

The GPU optimisation required a major redesign of the code which allowed to avoid transpositions of small temporary arrays. This optimisation can be applied as well to the CPU version. As a first step we applied it to a serial version of the spectral transform on CPU which is now used operationally at ECMWF and provides a major speedup as shown in the paper. Applying the same idea to the parallel CPU code requires more work that is planned in the future.

There is of course some risk that different optimisations are conflicting. We have not encountered this issue in ESCAPE so far. In the end we take whichever optimisation gives us the best performance in the full model. We currently have different source codes for CPUs and GPUs. We try to achieve a single source code through the use of the DSL except for highly specialised low level libraries like the spectral transforms.

Figure 11, the "barrier" optimisation is not clearly explained. Again, line plots of noncontinuous variable should be changed.

In the original code, some MPI barriers were present to profile the MPI communications at low level. By disabling these barriers (as allowed and documented in the code), a performance gain was identified. The main lesson learnt by this removal is the following: these barriers imply useless synchronisations which have the following consequences:

1. The application suffers from imbalance (as reported in D3.3 section 4.3.1.3 on a single node run). Thus, adding non mandatory synchronisation via barriers decreased the global performance (as each barrier implies to wait for the slowest process);
2. Moreover, these synchronisations created contention;

3. Last but not least, due to the two first points which change the behaviour of the application, there is a bias in the communications profiling. In other words, this profiling change the application behaviour.

As in the response to the comment on figure 10 we made the line thinner and added a statement to the caption.

Response to the referee 3 for “The ESCAPE project: Energy-efficient Scalable Algorithms for Weather Prediction at Exascale”

(grey background: text of the reviewer comment, white background: our response)

General comment:

In this manuscript an overview about the achievements in the ESCAPE project is given. The main concept is explained, some of the developments are explained in details and finally some tests are mentioned.

Although I think that this manuscript is a valuable contribution for GMD, I cannot recommend to accept the manuscript in the actual state. The manuscript must be revised in a substantial way before it can be considered again. Therefore I recommend major revision of the manuscript. In the following I will explain my concerns.

Major issues

1. Balance of the manuscript: The manuscript is very long and not really balanced. Some parts are explained in details, as e.g. the development of the MPDATA dwarf, but some parts are just mentioned. Especially for the very shortly explained parts, there are very often references to technical reports, i.e. documentation which is generally not peer reviewed. Although there are some performance tests, there is only one figure showing a test for atmospheric flows, and also this test is only marginally described.

The purpose of the paper is to present the concept of the dwarfs and to describe our work with a few detailed examples. Following the suggestions by reviewers 1 and 2 we have significantly shortened the paper. Having all of the details from the technical reports inside this paper would make it too long and would distract from the main message of the paper. To our knowledge GMD allows references to non peer reviewed technical reports if no peer reviewed reference is available.

The optimisations performed in the ESCAPE project do not affect the accuracy of the result. The figure which showed a comparison between finite volume and spectral transform method was an additional information. We removed this figure in our effort to shorten the paper.

I would recommend to significantly reorganize the manuscript, maybe also considering to split the manuscript into three parts: First, an overview part, where mostly the concept and the new architecture can be explained in a concise way. Second, a model description part, i.e. a detailed description of the different parts of the model, especially of the parts, which are contained in the technical memoranda but not described in peer-reviewed literature. Third, a part dedicated to test cases for atmospheric flows - and maybe also clouds and radiation, since these parts are also included into the model.

The suggested structure of the paper does not fit to the intended purpose of the paper to introduce the dwarf concept and illustrate the workflow with a detailed example. According to the

advice from reviewer 1 and 2 we have significantly shortened the paper with the goal to make the intended message clearer. As stated before the results of test cases for atmospheric flows can be found in the literature given in the references. In this paper we allowed optimisations only if they did not affect the results in any significant way. Under this constraint a description of the optimisation and the performance analysis are sufficient to present our results.

Especially test cases of atmospheric flows would be very interesting, since it is not clear if all the new models represent the atmospheric flow and other atmospheric phenomena in a physically consistent way. Therefore I highly recommend to use well-documented test cases for atmospheric flows, as e.g. Jablonowski & Williamson (2006). It would be interesting to see also tests for clouds and radiation, although I am not really aware of large scale tests, beyond the standard tests as e.g. Weismann & Klemp (1982).

As stated before our work on optimising code does not affect the accuracy of the results. Test cases for the underlying methods can be found in the literature referenced throughout the paper.

2. Selection of the dwarfs: It is not really clear how and why the different dwarfs were chosen. Although I think that this is a well chosen sample of possible models, it should be justified much better. Especially, the choice of the shallow water model is not really clear, because no real results of this model are shown in the manuscript. Therefore, I recommend to describe the choice of the models in a clearer way.

We have revised our conclusions section to make it clearer that we intend this paper to be a starting point and with a hope that the community will join our efforts and identify characteristic patterns in terms of computation and communication (dwarfs) and implement prototypes which can be used to work on optimising the key building blocks of weather & climate models.

Minor issues:

Cost model: The benefit of the cost model is not really clear to me. It is introduced in a comparable length as the dwarfs, but it is not really clear why this is so important for the whole manuscript, justifying a large part in the appendix.

The cost model was meant to illustrate the importance of porting as much of the model to the accelerator as possible. We have removed the cost model in order to shorten the paper as requested by reviewers 1 and 2.

The ESCAPE project: Energy-efficient Scalable Algorithms for Weather Prediction at Exascale

Andreas Müller¹, Willem Deconinck¹, Christian Kühnlein¹, Gianmarco Mengaldo¹, Michael Lange¹, Nils Wedi¹, Peter Bauer¹, Piotr K. Smolarkiewicz¹, Michail Diamantakis¹, Sarah-Jane Lock¹, Mats Hamrud¹, Sami Saarinen¹, George Mozdzyński¹, Daniel Thiemert¹, Michael Glinton², Pierre Bénard², Fabrice Voitus², Charles Colavolpe², Philippe Marguinaud², Yongjun Zheng², Joris Van Bever³, Daan Degrauwe³, Geert Smet³, Piet Termonia^{3,4}, Kristian P. Nielsen⁵, Bent H. Sass⁵, Jacob W. Poulsen⁵, Per Berg⁵, Carlos Osuna⁶, Oliver Fuhrer⁶, Valentin Clement⁷, Michael Baldauf⁸, Mike Gillard⁹, Joanna Szmelter⁹, Enda O'Brien¹⁰, Alastair McKinstry¹⁰, Oisín Robinson¹⁰, Parijat Shukla¹⁰, Michael Lysaght¹⁰, Michał Kulczewski¹¹, Milosz Ciznicki¹¹, Wojciech Piątek¹¹, Sebastian Ciesielski¹¹, Marek Błażewicz¹¹, Krzysztof Kurowski¹¹, Marcin Procyk¹¹, Pawel Spychala¹¹, Bartosz Bosak¹¹, Zbigniew Piotrowski¹², Andrzej Wyszogrodzki¹², Erwan Raffin¹³, Cyril Mazauric¹³, David Guibert¹³, Louis Douriez¹³, Xavier Vigouroux¹³, Alan Gray¹⁴, Peter Messmer¹⁴, Alexander J. Macfaden¹⁵, and Nick New¹⁵

¹European Centre for Medium-Range Weather Forecasts (ECMWF), Reading, UK

²Centre National de Recherches Météorologiques, Météo-France, Toulouse, France

³Royal Meteorological Institute (RMI), Brussels, Belgium

⁴Department of Physics and Astronomy, Ghent University, Ghent, Belgium

⁵The Danish Meteorological Institute (DMI), Copenhagen, Denmark

⁶Federal Office of Meteorology and Climatology MeteoSwiss, Zurich, Switzerland

⁷Center for Climate System Modeling, Zurich, Switzerland

⁸Deutscher Wetterdienst (DWD), Offenbach, Germany

⁹Loughborough University, Leicestershire, UK

¹⁰Irish Centre for High-End Computing (ICHEC), National University of Ireland, Galway, Ireland

¹¹Poznan Supercomputing and Networking Center (PSNC), Poznań, Poland

¹²Institute of Meteorology and Water Management - National Research Institute (IMGW-PIB), Warsaw, Poland

¹³ATOS, Bezons, France

¹⁴NVIDIA Switzerland, Zurich, Switzerland

¹⁵Optalysys Ltd., Glasshoughton, UK

Correspondence: Andreas Müller (andreas.mueller@ecmwf.int)

Abstract. In the simulation of complex multi-scale ~~flow problems, such as those flows~~ arising in weather and climate modelling, one of the biggest challenges is to satisfy ~~operational-strict service~~ requirements in terms of time-to-solution and energy-to-solution~~yet,~~ without compromising the accuracy and stability of the calculation. These competing factors require the development of ~~state-of-the-art robust~~ algorithms that can optimally exploit the targeted underlying hardware and efficiently deliver the extreme computational ~~capabilities efforts~~ typically required in operational forecast production. These algorithms should: (i) minimise the energy footprint along with the time required to produce a solution~~;~~; (ii) maintain ~~a satisfying the scientifically required~~ level of accuracy~~;~~; and (iii) be numerically stable~~and resilient,~~ ~~and resilient~~ in case of hardware or software failure.

The European Centre for Medium Range Weather Forecasts (ECMWF) is leading a project called ESCAPE (Energy-efficient Scalable Algorithms for weather Prediction on Exascale supercomputers) which is funded by Horizon 2020 (H2020) under initiative Future and Emerging Technologies in High Performance Computing (FET-HPC). The goal of the ESCAPE project is to develop a sustainable strategy to evolve weather and climate prediction models to next-generation computing technologies. The project partners incorporate the expertise of leading European regional forecasting consortia, university research, experienced high-performance computing centres and hardware vendors.

This paper presents an overview of ~~results obtained in the ESCAPE project in which weather prediction have been broken down into smaller building blocks called dwarfs~~ this ESCAPE strategy: (i) identify domain-specific key algorithmic motifs in weather prediction and climate models (Weather & Climate dwarfs), (ii) categorise them in terms of computational and communication patterns, while (iii) adapting them to different hardware architectures with alternative programming models, (iv) analyse the challenges in optimising and (v) finding alternative algorithms for the same scheme. The participating weather prediction models are: IFS (Integrated Forecasting System), ~~ALARO~~, ALARO, a combination of AROME (Application de la Recherche à l'Opérationnel à Meso-Echelle) and ALADIN (Aire Limitée Adaptation Dynamique Développement International); and COSMO-EULAG, a combination of COSMO (Consortium for Small-scale Modeling) and EULAG (Eulerian/semi-Lagrangian fluid solver). ~~The dwarfs are analysed and optimised in terms of computing performance for~~ For many of the Weather & Climate dwarfs ESCAPE provides prototype implementations on different hardware architectures (mainly Intel Skylake CPUs, NVIDIA GPUs, Intel Xeon Phi). ~~The ESCAPE project includes the development of new algorithms that are specifically designed for better energy efficiency and improved portability through domain specific languages. In addition, the modularity of the algorithmic framework, naturally allows testing different existing numerical approaches, and their interplay with the emerging heterogeneous hardware landscape. Throughout the paper, we will compare different numerical techniques to solve the main building blocks that constitute weather models, in terms of energy efficiency and performance, on a variety of computing technologies.~~ with different programming models. The spectral transform dwarf represents a detailed example of the co-design cycle of an ESCAPE dwarf.

1 Introduction

Numerical weather prediction (NWP) has made significant progress over the and climate prediction capabilities represent substantial socio-economic value in nearly all sectors of human society, namely for the mitigation of the impact of extremes, in food production, renewable energy and water management, infrastructure planning, and for finance and insurance where weather sensitive goods and services are traded. Despite significant progress achieved over past decades (Bauer et al., 2015) ~~Nevertheless, the prediction of extreme weather events, like heavy precipitation patterns which could lead to flooding and tropical cyclones, is still far from being satisfying. An improvement in the forecast of these events would have a huge economic and societal benefit.~~ there are substantial shortcomings in our ability to predict, for example, weather extremes with sufficient lead time and the impact of climate change at regional or national level. Extreme weather events caused over 500 thousand casualties and over 2 trillion \$US economic damages in the past 20 years (Wallemaq et al., 2018). Another example is the

~~prediction of climate change mitigation and adaptation targets. Failure to meet these targets is ranked among the leading threats to global society (World Economic Forum, 2019).~~

~~One of the main sources of errors in NWP is insufficient resolution. Global NWP and climate simulations for the foreseeable future will not reach what may be considered “sufficient resolution”, which for practical NWP purposes may be in the key~~
5 ~~sources of model error is limited spatial (and temporal) resolution (Palmer, 2014) which implies that key physical processes that drive global circulation, like deep convection in the tropics and mesoscale eddies in the ocean, are only crudely represented in models by so-called parameterisations. In addition, deficiencies in the representation of process interactions between atmosphere and ocean/sea-ice, as well as atmosphere and land, including a time-varying biosphere, are highly relevant strategic targets to improve the representation of the internal variability of the Earth system in models (Katzav and Parker, 2015). However,~~
10 ~~better spatial resolution and enhanced model complexity translate immediately into significant computational challenges (Schulthess et al., 2019), whereby spatial resolution is the most demanding because a doubling of resolution roughly translates into eight times more computations (doubling the horizontal resolution in both directions and a corresponding decrease in the time step). The critical resolution threshold at which global weather & climate models may eventually be able to overcome the bulk of the limiting deficiencies is unclear; however, O(1km global resolution range. As a result a substantial part of the cost~~
15 ~~of a model timestep is used in physical parameterisations that describe the effect of unresolved processes on the resolved scale, but also to describe diabatic effects such as radiation and water phase changes.~~

~~One of the key components of the ESCAPE project is to guarantee the continued efficiency of NWP models while transitioning to emerging computing architectures including accelerators. A particular challenge arises from the need to achieve time-to-solution and energy-to-solution efficient solutions for operating global, complex, high-resolution, ensemble based systems on high-performance~~
20 ~~computers so that they will remain affordable given tight operational schedules. A comprehensive assessment of the modelling infrastructure of ECMWF’s Integrated Forecasting System (IFS) can be found in Wedi et al. (2015). The authors there also stress a need to enhance model complexity where it improves forecast skill through fully coupled simulations of the atmosphere with ocean, sea-ice and land surfaces and including interactive chemical processes in the atmosphere) emerges as a defensible intermediate target (Shukla et al., 2010; Neumann et al., 2019).~~

~~Beside The ESCAPE project has been motivated by the need for increased model complexity we also see a need for~~
25 ~~significantly increased resolution. It is ECMWF’s strategic goal by 2025 to establish the efficiency required to run at least twice daily global ensembles at 5km globally uniform resolution. The potential scope of such an ensemble is illustrated best in the context of extreme events such as tropical storm Irma. Experiments with a 5km resolution ensemble show dramatic improvement in the forecast of running Earth-system models at much higher resolution and complexity than presently available, but within the same time-critical path of daily production for weather forecasts and with the same production throughput~~
30 ~~needed for decadal/centennial climate projections as used today. This translates to computing one simulated year per wallclock day. Energy efficiency is a key requirement as power envelopes for HPC systems cannot be scaled up at the same rate as the computing demand. Obviously, this presents a substantial challenge to all aspects of computing, namely the choice and implementation of numerical methods and algorithms, and the programming models to map memory access and communication patterns onto specific hardware (Wehner et al., 2011).~~

75 The ECMWF specific strategic target is to provide skilful predictions of high-impact weather up to two weeks ahead by 2025 based on a global 5 km-scale Earth-system model ensemble, complemented by the corresponding, high-resolution data assimilation system for creating the initial conditions. This system needs sufficient model complexity in the atmosphere (physics and chemistry), oceans, sea-ice and land to represent all processes acting on such scales, and a sufficiently large ensemble size so that complex probability distributions of extreme forecast features can be sampled well enough. In terms of computational demand this is comparable to the ~~wind intensity compared to the currently operational resolution of 18km~~ (Magnusson et al., 2018).

—ECMWF is world leading in terms of track forecast accuracy and increasing the resolution to 5km does not compromise the accuracy of the track forecast.

above 1-km resolution target for a single forecast.

85 The supercomputers used for ~~NWP~~ numerical weather prediction (NWP) have changed dramatically over the past decades, and ECMWF's Integrated Forecasting System (IFS; Wedi et al., 2015, and references therein) has exhibited rather different levels of sustained performance ranging from 40-50% on the parallel vector-processor Cray-XMP, YMP and C90 architectures between 1984 and 1996, to 30% on Fujitsu's similar-type VPP-700 and 5000 systems between 1996 and 2002, down to 5-10% on the scalar multi-core IBM P5-7 and Cray XC30 and 40 architectures operated between 2002 and today. Despite ~~sustained performance declining~~, overall FLOP performance increased exponentially thanks to Moore's law, Dennard scaling and processor pricing (e.g. Flamm, 2018) so that significant upgrades of model resolution, complexity and ensemble size remained affordable in terms of capital cost and power supply. As this 'natural' technological evolution is slowing down, new concepts for designing algorithms and for mapping the associated computational patterns onto the available architectures - from processor to system level - are needed (Lawrence et al., 2018). Many of the algorithms used in NWP were designed well ~~before the multi-core era started~~. ~~Optimisations have been continuously developed with every new computer that was bought by NWP centres. These optimisations were usually limited to adapting the same FORTRAN code base.~~

The , but even though they contain highly-tuned shared and distributed memory parallelisation they only achieve such limited sustained performance because of poor arithmetic density (ratio of data communication / computations) in some of the highly varying kernels, sequential tasking due to strong scientific dependencies, and rather complex algorithms required to ~~solve~~ a multi-scale and multi-phase fluid dynamics problem on a rotating sphere with sufficient accuracy.

If the envisioned increase in ~~the resolution of NWP models together with the increased model complexity~~ requires a ~~much~~ model fidelity is constrained by only marginally growing power envelopes and decelerating general purpose processor speed, then performance issues need to be addressed at the root, and a more radical redesign of the basic algorithms and codes used for weather prediction ~~needs~~ to be considered. This is why ESCAPE investigates both HPC adaptation and ~~alternative~~ numerical formulations in scientifically and computationally well defined model components, the Weather & Climate dwarfs, followed by a detailed analysis of their respective computational bottlenecks, and subsequent hardware and algorithm dependent optimisations.

We start with introducing the concept of The Weather & Climate Dwarfs dwarf idea is introduced in Section 2. In Section ?? we present our work on optimising codes used in the dynamical cores and Section ?? shows results for physics

parameterisations we illustrate the usefulness of the dwarf concept with the example of the spectral transform dwarf. The paper ends with a comparison between different methods in Section ?? and conclusions and outlook in Section 4.

2 The dwarf concept Weather & Climate dwarfs

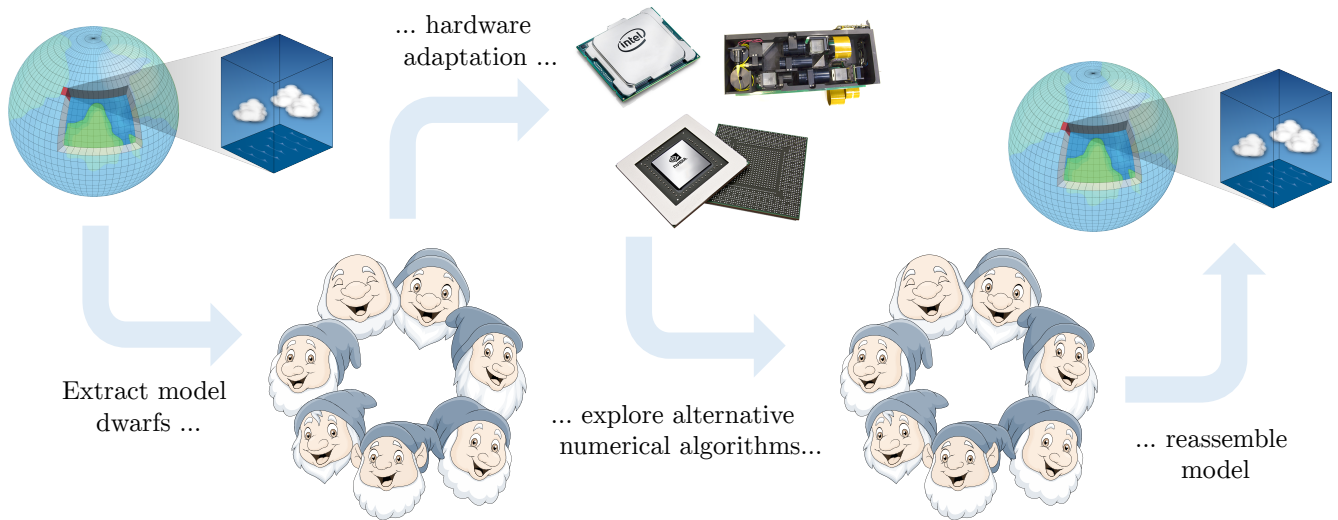
5 Weather prediction and climate models are too big to adapt to new computing architectures at once. For this reason we disassemble the model into smaller building blocks which we call Weather & Climate Dwarfs (Figure 1), in analogy to the Berkeley Dwarfs (Asanović et al., 2006). We define a Weather & Climate Dwarf as a key functional component of an earth system model. Each dwarf is associated with specific

2.1 Motivation

10 In 2004 Phillip Colella introduced the seven dwarfs of algorithms for high-end simulation in the physical sciences (Colella, 2004). These were later extended to the 13 Berkeley dwarfs (Asanović et al., 2006, 2009) which are meant to represent characteristic patterns of computation and communication patterns which lead to characteristic bottlenecks relevant for the performance of. These dwarfs were created to cover the characteristic computational properties of a broad range of scientific applications. These dwarfs are the basis for the OpenDwarfs benchmark suite (Feng et al., 2012; Krommydas et al., 2015; Johnston and Milthorpe, 2018) and were also applied to benchmark cloud computing in Phillips et al. (2011). In a similar fashion Kaltofen (2011) introduced the
15 seven dwarfs of symbolic computation.

Following this idea, we categorise key algorithmic motives specific to weather prediction and climate models and identify their specific computational and communication patterns, which in return are crucial for the overall model performance. In entire model performance. The dwarfs thus represent domain specific mini-applications (Messer et al., 2016) which include direct input from the domain scientist, documentation, timers for profiling purposes as well as error estimates for verification purposes. In
20 this way the dwarfs facilitate communication of weather domain specific knowledge and algorithmic motifs with specialists in other domains. Different implementations of the dwarfs can be used as a first step towards optimising and adapting weather prediction and climate models to new hardware architectures and to benchmark current and future supercomputers with these simple but relevant applications. Identifying these key algorithms also allows better collaboration between operational weather prediction centres, hardware vendors and academia because they are much simpler to understand than the full models. The
25 concept of dwarfs is different from the existing separation of weather and climate models into different model components, such as atmosphere, land-surface and ocean for which separate dynamical core and physical parameterisation packages already exist. Instead, dwarfs define a runnable and more manageable sub-component in a hierarchy of model complexity for specific targets such as adaptation to GPUs, exploring alternative programming models, and developing performance portable domain specific languages. But dwarfs can also be used by domain scientists for developing alternative algorithms.

30 The fundamental starting point of the ESCAPE project we adapt these dwarfs to different hardware architectures. The knowledge gained in this adaptation is used to research alternative numerical algorithms which are better suited for those new architectures. Eventually the goal is to speedup the entire simulation by using the optimised dwarfs in the full model ESCAPE



Earth illustration: used under license from GraphicsRF/Shutterstock.com.
 Dwarf illustrations: used under license from Teguh Mujiono/Shutterstock.com

Figure 1. Illustration of the main idea behind the ESCAPE project. The entire model is broken down into smaller building blocks called dwarfs. These are adapted to different hardware architectures. Based on the feedback from hardware vendors and high performance computing centres alternative numerical algorithms are explored. These improvements are eventually built into the operational model.

project is to identify the dwarfs in the participating weather and climate models (Figure 1) and to adapt them to different hardware architectures. The knowledge gained in this adaptation is used to research alternative numerical algorithms which are better suited for those new architectures, and experiment with alternative programming models, towards improving the overall energy-efficiency of weather prediction applications.

5 2.2 The data structure framework Atlas

To avoid code duplication across dwarfs, we use a data structure framework called Atlas, that handles both the mesh generation and parallel communication aspects. The main idea behind Atlas (library developed at ECMWF, Deconinck et al., 2017) is illustrated in Figure ??a : Atlas provides functionality to generate the grid, create the mesh and partition it. The partitioning step also generates halos (overlapping regions between different MPI partitions to enable stencil operations in a parallel context), if needed by the numerical method. Atlas provides *function spaces* for different numerical methods as well as storage objects called *fields*, which can be used by the dwarfs to store their data. More information about Atlas can be found in Deconinck et al. (2017). (a) Illustration of the concept behind the data structure framework Atlas. Further explanation can be found in the main text. (b) Example of the limited-area extension of Atlas. Shown here is a Lambert projection centred around 40longitude and 50latitude. The grid consists of 50 grid cells in zonal direction, 40 grid cells in meridional direction and the resolution is in both directions 0.5. Color shading shows the domain decomposition.

In the ESCAPE project, Atlas has been improved and extended by adding support for limited area grids (Figure ??b) and adding support for accelerators through GridTools (Deconinck, 2017a, b).

2.2 List of dwarfs created in ESCAPE

Table 1 gives an overview of the dwarfs defined in the ESCAPE project. These dwarfs have been chosen because they have a significant contribution to the overall runtime of many represent key algorithms that are representative of any weather and climate model, or because their computational and communication patterns represent one of the most runtime consuming parts of weather forecasting systems (Figure 2). In addition they represent fundamental computational patterns that are also relevant for those parts of the model that are not explicitly covered by these dwarfs. For many of the dwarfs we created so called prototypes. Each prototype implementation for a specific hardware (by default CPU) comes with documentation (Mengaldo, 2016; Müller et al., 2017), error measures which allow us to quickly see if optimisations affect the results and timers to see the speedup obtained through optimisations. Each of the released dwarfs has been documented (Mengaldo, 2016; Müller et al., 2017). Table 1 also lists the models from which the dwarf code in which we identified the dwarf and from which the prototype implementations originated. The models include the spectral IFS, the IFS-finite volume module IFS with the spectral transform method (IFS-ST), the Finite-Volume Module of the IFS (IFS-FVM), ALARO/AROME and COSMO-EULAG. As part of the ESCAPE project In order to explore alternative discretizations in ESCAPE, we created a new global shallow water model called GRASS (Global Reduced A-grid Spherical-coordinate System). This model is introduced in Section ??.

Table 1 further shows which dwarfs are using Atlas for handling the mesh. This table have prototypes that are based on the data structure framework Atlas. We use Atlas to ease adaptation to future architectures and avoid code duplication across prototypes. Atlas handles both the mesh generation and parallel communication aspects (Deconinck et al., 2017). In the ESCAPE project, Atlas has been improved and extended by adding support for limited area grids and adding support for accelerators through GridTools (Deconinck, 2017a, b). Table 1 also shows the programming model adopted for each dwarf of the dwarf prototypes. In particular, we used MPI for distributed-memory parallelism, OpenMP and OpenACC for shared-memory parallelism, and DSL, that stands for domain-specific language and uses GridTools. Also note that the spectral transform on the sphere with spherical harmonics and in limited area domains with biFFT has also been implemented with the optical processor by Optalysys (Maefaden et al., 2017). We will give an overview of the work with the optical processor in Section 3. Overview of work performed in the ESCAPE project, the corresponding Section in this paper and the models from which the dwarfs originated. The column Atlas shows which of the dwarfs are based on the data structure framework Atlas (see Section ??). Dwarfs for which MPI and OpenMP/OpenACC is available both can be used together as hybrid parallelisation. LAITRI is used in the full model in each MPI process individually and therefore does not include MPI parallelisation inside the dwarf. Further explanations can be found in the text.

In addition to identifying computational and communication patterns in existing models we also performed research on new algorithms and approaches which have the potential to reduce runtime and energy consumption significantly in the future. We developed a multigrid preconditioner for the Krylov-subspace solver employed in the semi-implicit time integration in IFS-FVM (Müller et al., 2017) to reduce the number of iterations in iterative solves, a HEVI time-integration scheme with




Figure 2. Portion of the forecast runtime spent in ESCAPE dwarfs for the three different models IFS (left), ALARO ensemble prediction system (EPS) (middle) and COSMO-EULAG (right). The measurements for IFS were taken during an operational run on 352 nodes (1408 MPI processes, 18 OpenMP threads per process). The limited area models ALARO-EPS and COSMO-EULAG used each 576 MPI processes for the simulations shown here. IFS and ALARO-EPS are both based on the spectral transform method with latter employing 2D FFTs. For this reason the two left pie charts share one legend. COSMO-EULAG uses different methods and the pie chart has its own legend. The vertical line separates the two different legends.

significantly improved stability (Colavolpe et al., 2017) to avoid some of the global communication patterns, and connected to this, explored alternative finite difference methods on the sphere (Bénard and Glinton, 2019; Glinton and Bénard, 2019). Finally, we explored FFTs and spherical harmonics on optical (co-)processors (Macfaden et al., 2017) to potentially scale these transforms towards higher resolutions at a fixed energy cost. An overview of the work with the optical processor can be found in section 3.5.

The work on the hardware adaptation and performance optimisation focused mainly on three of our dwarfs: the spherical harmonics spectral transform dwarf, the MPDATA advection dwarf and the ACRANEB2 radiation dwarf. The work on the radiation dwarf has already been described in detail in Poulsen and Berg (2017). We will give a short summary of that work in Section ?? and focus our description in Section ?? on the work on the spectral transform and MPDATA advection which has not been published elsewhere. For the same reason we will also describe the algorithms and computational challenges for these two dwarfs in greater detail.

In NWP, we distinguish between different components of our weather prediction systems. One of the main components is data assimilation which computes the initial conditions by using observations and previous forecasts. We call the component that takes this initial condition and computes the future weather the forecast component. As mentioned before, the dwarfs in

Table 1. Overview of dwarfs identified in the ESCAPE project, short description of the computational characteristics of each dwarf and the models in which we identified the dwarf. Standalone applications including documentation, profiling and verification capabilities called prototypes have been released for many of the dwarfs. The column Atlas shows which of the dwarfs have prototypes that are based on the data structure framework Atlas (Decoinek et al., 2017). Dwarfs for which prototypes with MPI and OpenMP/OpenACC is available can be used with both programming models together as hybrid parallelisation. LAITRI is used in the full model in each MPI process individually and therefore does not include MPI parallelisation inside the dwarf. Further explanations can be found in the references given in the last column.

	computational characteristics	IFS-ST	IFS-FVM	ALARO/AROME	COSMO-EULAG	GRASS	prototype released	Atlas	MPI	OpenMP	OpenACC	DSL (Section 3.8)	Optalysys (Section 3.5)	references
Dynamics dwarfs:														
• time-stepping:														
spectral transform:														
– spherical harmonics	matrix multiplication, global transposition	✓					✓	✓	✓	✓	✓	✓	✓	Mengaldo (2016)
– biFFT	1D and 2D FFT algorithm, global transposition iteration, global reduction		✓				✓	✓	✓	✓	✓	✓	✓	Mengaldo (2016)
elliptic solver - GCR	interpolation for remapping		✓				✓	✓	✓	✓	✓	✓	✓	Mengaldo (2016)
elliptic solver - multigrid	nearest neighbour column communication						✓ ¹	✓	✓	✓	✓	✓	✓	Müller et al. (2017)
HEVI								✓	✓	✓	✓	✓	✓	Mengaldo et al. (2018); Colavolpe et al. (2017)
• advection:														
semi-Lagrangian	large halo communication	✓		✓			✓	✓	✓	✓	✓	✓	✓	Mengaldo (2016)
– LAITRI (3d interpolation)	efficient data access	✓		✓			✓	✓	✓	✓	✓	✓	✓	Robinson et al. (2016)
MPDATA	small halo communication, indirect addressing ²		✓				✓	✓	✓	✓	✓	✓	✓	Mazaauric et al. (2017b); Osuna (2018)
• gradient computation:														
high order finite difference	large halo communication					✓		✓	✓	✓	✓	✓	✓	Bénard and Glimton (2019); Glimton and Bénard (2019)
finite volume	small halo communication, indirect addressing ²		✓		✓		✓	✓	✓	✓	✓	✓	✓	Kühnlein et al. (2019)
Physics dwarfs:														
• radiation scheme:														
ACRANE2	many transcendental functions			✓			✓	✓	✓	✓	✓	✓	✓	Müller et al. (2017); Poulsen and Berg (2017)
• cloud microphysics:														
CLOUDSC	single (vertical) column, many if-statements, many local variables	✓	✓				✓	✓	✓	✓	✓	✓	✓	Xiao et al. (2017); Clement et al. (2018)

¹ The GRASS model is currently only solving the shallow-water equations. Research on HEVI time integration methods was performed to explore future options.

² when using unstructured meshes

Table 1 have a significant contribution to the overall runtime of the forecast component. The ESCAPE project has not explicitly covered code used for data assimilation. However, several algorithmic motifs found in data assimilation can be found in the collection of dwarfs. In addition the cost of data assimilation depends very much on the method that is used. Some NWP centres use 3DVAR like the new ensemble prediction system ensemble prediction system developed at RMI in Belgium (RMI-EPS).

5 Our studies in ESCAPE have shown that for RMI-EPS, 99% of the energy consumption is spent in the forecast component. Data assimilation takes less than 1% of the energy consumption even though it uses 35% of the runtime. The reason for this is that most computations inside the data assimilation in RMI-EPS are running on a small number of nodes which gives them a relatively large contribution in terms of runtime and very small contribution in terms of energy consumption. More details about this study of RMI-EPS can be found in Van Bever et al. (2018b). The 4DVAR data assimilation used at ECMWF is much

10 more expensive but most of the runtime of 4DVAR is spent in linearised versions of the forecast model (Hamrud, 2010). The dwarfs represent therefore also significant components of As an example of the cycle involved in identifying, dwarfing, testing, adapting, optimising, and considering alternative solution procedures for a specific algorithmic motif in the ESCAPE project, we shall illustrate in the next section the work on the 4DVAR data assimilation at ECMWF, spectral transform dwarf in detail.

2.3 Simple theoretical model for using accelerators

15 Before we describe how the different dwarfs have been adapted and optimised for different hardware architectures, we first discuss the general benefit of using hybrid machines with CPUs and accelerators in a simple cost model. It is well known that data transfer between CPUs and accelerators on hybrid machines can take a significant amount of time and needs to be avoided as much as possible. Data transfer can be avoided to some degree by running the entire model on the accelerator and transferring only the initial condition at the beginning of the simulation to the accelerator and the necessary output back to the

20 CPU when required. In our simple cost model we show that even if we neglect the overhead caused by data transfer we still get a significant overhead if the computation on the CPU cannot overlap with computation on the accelerator. Unfortunately many NWP models are currently not able to overlap these computations.

3 Dwarf example: spectral transform

The derivation of our simple cost model can be found in the Appendix ???. This cost model considers specifically the case of

25 the NWP centres where the runtime is fixed and the number of processors need to be adjusted to reach the given runtime. If the simulation started earlier we would not have all of the necessary observations and if the simulation finished later the result would not reach our customers on time. To get some practical numbers from our theoretical model we use the situation at Météo-France as an example. If we We start with a CPU-only code and port a small part of the code to the accelerator the total cost per time-step first increases if we cannot overlap computations on the CPU with computations on the accelerator (solid

30 lines in Figure ??a). This overhead can be avoided if we find a way to overlap computations on the CPU with computations on the accelerator (dashed lines in Figure ??a). Impact of (a) speedup of the accelerator and (b) strong scaling efficiency on the total cost of one time-step as a function of how much of the code is running on the accelerator. This Figure is based on the

specific case of the AROME model at Météo-France. The strong scaling efficiency of 53% (used for all curves in (a) and for the red curve in (b)) at 1000 CPU sockets compared to one CPU socket was found through linear regression experimentally and should therefore be the most realistic case for the AROME model. Dotted lines illustrate the result for two of the curves if we assume perfect overlap between computations on the CPUs and accelerators while the solid lines assume no overlap. In

5 (b) we assume an accelerator speedup of 8x. short description of the (domain-specific) spectral transform and continue with a subsection discussing computational challenges with a particular focus on the data structures and data access patterns used in IFS. After that we present the work that has been done on adapting and optimising the dwarf for GPUs, CPUs and optical processors. We finish with a comparison between the results obtained for the different architectures and a discussion on the sustainability of the chosen techniques.

10 The sensitivity of the results on the scaling efficiency of the model is shown in Figure ??b. The cost decreases with increasing strong scaling efficiency. Our simple cost model assumes that CPU and accelerator part of the code have the same strong scaling efficiency. Under this assumption fast accelerators would have a larger benefit for code that has worse scaling efficiency because the speedup of the accelerator would allow to reduce the number of devices. Having the same scaling efficiency on CPU and accelerator will be difficult to achieve in reality because accelerators usually require a lot of inherent parallelism inside each

15 device.

We conclude from this simple cost model that having no overlap between computations on the CPU and on the accelerator can produce a significant overhead. This overhead can be minimised by running most of the NWP model on the accelerator (as done in Fuhrer et al., 2018). At the same time we should invest in research on overlapping computations on CPUs and accelerators. Overlapping these computations is not part of the work done in the ESCAPE project. We plan to investigate this

20 overlapping in our future research.

In the next two sections, we introduce the dynamics dwarfs (Section ??) and the dwarfs associated to the physical parametrisations (Section ??). This conceptual categorisation follows the rationale adopted for classifying the various dwarfs mentioned above and reported in Table 1.

4 Dynamics dwarfs

25 This section describes the work done in the ESCAPE project on the dynamics dwarfs which are used in the dynamical cores of the models. The subsections follow the list of dynamics dwarfs as given in Table 1. Three important features shared by these dwarfs are: a.) solution of the governing equations in the classical meteorological latitude-longitude coordinates, b.) compatibility with reduced Gaussian grids, c.) co-located (i.e. unstaggered) arrangement of the prognostic variables with respect to the grid. Having these shared features between the dwarfs facilitates incorporation into the IFS code infrastructure.

3.1 Spectral transform Background

3.1.1 Background

Each time-step of IFS is split between computations in ~~grid point space and computations in spectral space~~ physical space (i.e. a grid point representation) and computations with respect to spectral space (i.e. spectral coefficients at different wavenumbers).

5 Semi-Lagrangian advection, physical parameterisations and products of terms are computed most efficiently in grid point space while horizontal gradients, semi-implicit calculations and horizontal diffusion are computed more efficiently in spectral space. The transform between these two spaces is performed on the sphere with spherical harmonics, that is computing these results along longitudes in a Fast Fourier transform (FFT) and a Legendre transform (LT) along latitudes. Limited area models replace the Legendre transform with another FFT which leads to the name biFFT.

10 In spectral transform methods such as the one used in IFS (Wedi et al., 2013), the specific form of the semi-implicit system facilitating large time-steps (and thus time-to-solution efficiency) is derived from subtracting a system of equations, linearised around a horizontally homogeneous reference state. The solution of this linear system is greatly accelerated by the separation of the horizontal and the vertical part, which matches the large anisotropy of horizontal to vertical grid dimensions prevalent in atmospheric and oceanic models. In spectral transform methods one uses the special property of the horizontal Laplacian
15 operator in spectral space on the sphere

$$\nabla^2 \psi_n^m = -\frac{n(n+1)}{a^2} \psi_n^m, \quad (1)$$

where ψ symbolises a prognostic variable, a is the Earth radius, and (n, m) are the total and zonal wavenumbers of the spectral discretisation (Wedi et al., 2013). This conveniently transforms the 3D Helmholtz problem into an array (for each zonal wavenumber) of 2D matrix operator inversions with the dimension of the vertical levels square, or in the case of treating
20 the Coriolis term implicitly, vertical levels times the maximum truncation, resulting in a very cheap direct solve.

In this paper we focus on the computational aspects and especially on the data layout. We illustrate here the inverse spectral transform on the sphere which goes from spectral space to grid point space. The direct transform adds one numerical integration. Otherwise ~~it works in the same way but in opposite direction~~ the only change between inverse and direct transform is that the direct transform starts with the Fourier transformation and applies the Legendre transformation afterwards.

25 The inverse spectral transform begins with the spectral data $\mathbf{D}(f, i, n, m)$ which is a function of field index f (for the ~~variables pressure~~ variable surface pressure at a single level and for wind vorticity, wind divergence and temperature at each height level), real and imaginary part i and wave numbers (zonal wave number $m = 0, \dots, N_T$ and total wave number $n = 0, \dots, N_T - m$ where N is the spectral truncation). Please note that we deviate here from the usual notation where total wavenumber goes from m to N_T because this simplifies the separation between even and odd n . We use here column-major order like in
30 ~~Fortran~~ FORTRAN, i.e. the field index f is the fastest moving index and the zonal wave number m is the slowest moving index. Typical dimensions can be seen in the operational high resolution (9km) forecast run at ECMWF: the number of fields is in this case 412 ~~and the~~ for the direct transform and 688 for the inverse transform and the number of zonal wave numbers is given by

the truncation $N_T = 1279$. The number of latitudes is $2N_T + 2 = 2560$ and the number of longitudes increases linearly from 20 next to the poles to $4N_T + 20 = 5136$ next to the equator.

We take advantage of the symmetry of the Legendre polynomials for even n and anti-symmetry for odd n . The coefficients of the Legendre polynomials are pre-computed and stored in $\mathbf{P}_{e,m}(n, \phi)$ for even n and $\mathbf{P}_{o,m}(n, \phi)$ for odd n , where ϕ stands for the latitudes of our Gaussian mesh. Only latitudes on the northern hemisphere are computed. Latitudes on the southern hemisphere are reconstructed from the northern latitudes as we will show later. In the same way we split the spectral data for each m into even part $\mathbf{D}_{e,m}(f, i, n)$ and odd part $\mathbf{D}_{o,m}(f, i, n)$. We write variables over which we can parallelise our computations as indices. The inverse Legendre transform is performed by computing the following matrix multiplications using BLAS:

$$\begin{aligned} \mathbf{S}_m(f, i, \phi) &= \sum_n \mathbf{D}_{e,m}(f, i, n) \cdot \mathbf{P}_{e,m}(n, \phi), \\ \mathbf{A}_m(f, i, \phi) &= \sum_n \mathbf{D}_{o,m}(f, i, n) \cdot \mathbf{P}_{o,m}(n, \phi). \end{aligned} \quad (2)$$

The resulting array for the symmetric and antisymmetric parts are now combined into the Fourier coefficients on the northern and southern hemisphere:

$$\begin{aligned} \phi > 0 : \mathbf{F}(i, m, \phi, f) &= \mathbf{S}_m(f, i, \phi) + \mathbf{A}_m(f, i, \phi), \\ \phi < 0 : \mathbf{F}(i, m, \phi, f) &= \mathbf{S}_m(f, i, \phi) - \mathbf{A}_m(f, i, \phi). \end{aligned}$$

These Fourier coefficients are finally used to compute the fields in grid point space at each longitude λ via FFT:

$$\mathbf{G}_{\phi, f}(\lambda) = \text{FFT}(\mathbf{F}_{\phi, f}(i, m)). \quad (3)$$

3.1.1 Computational challenges

3.2 Computational challenges

The computations in grid point space and spectral space require all the fields f to be on the same computational node. The summation over the total wavenumber n in the Legendre transform (2) makes it most efficient to have all total wavenumbers on the same node and the Fourier transform (3) over (i, m) makes it most efficient to have all of the zonal wavenumbers m with real and imaginary part on the same node. This is only possible if the data is transposed before and after the spectral transform as well as in between Legendre and Fourier transform. These transpositions produce substantial communication which increases the contribution of the spectral transform to the overall runtime for future resolutions (Figure 3).

Simplified simulations of the MPI communications performed in ESCAPE show that the strong scalability of the communication time for the spectral transform transpositions is better than for halo communication required by semi-Lagrangian advection and global norm computation commonly used in semi-implicit methods (Figure ??). ~~Simulation of the strong scaling of MPI communication time for minimal routing algorithm and dragonfly topology adapted from Zheng and Marguinaud (2018). The result for the spectral transform ends at 2×10^5 processes because larger number of processes exceeded the maximum~~




Figure 3. Cost profiles of the significant components of the IFS NWP model in percent of CPU time, at the operational 9km horizontal resolution with 137 vertical levels (left), the anticipated future horizontal resolution of 5km with 137 vertical levels (middle) and an experimental resolution of 1.25-1.45km with 62 vertical levels (right). The [GP-grid point](#) dynamics represents the advection and gridpoint computations related to the dynamical core, [SI-semi-implicit](#) solver represents the computations and communications internal to spectral space, [SP-transforms](#) spectral transform relates to the communications and computations in the transpositions from gridpoint to spectral and reverse as well as the FFT and DGEMM computations (see also spectral transform schematic below), [PHYSICSphysics+RAD-radiation](#) relates to the cost of physical parametrisations including radiation, and finally accounting for the additional components of the wave and the ocean model. The simulation at 1.25-1.45km is without ocean and waves. [All of these profiles have been obtained through measurements on the Cray XC40 supercomputer of ECMWF.](#)

runtime of the cluster used for these simulations of 24 hours. Halo communication for the semi-Lagrangian advection is assumed to use 20 elements which resembles the operational setting in IFS. The semi-implicit method uses MPI all-reduce to gather data from all processes and send the resulting norm to all of the processes. The reason for the relatively good scalability of the spectral transform is that the transpositions are not global but each transposition acts only on a much smaller communicator. The transposition between the computations in spectral space and the Legendre transform is exchanging field index f and total wavenumber n . This transposition is therefore independent between different zonal wavenumbers m . The transposition between Legendre and Fourier transform exchanges zonal wavenumber m with latitude ϕ and is independent between different fields. Finally the transposition between Fourier transform and computations in grid point space exchanges longitude λ and field index f and is independent between different latitudes ϕ . The transposition between Legendre and Fourier transform is therefore most costly because the number of independent communicators is with the number of fields much smaller than for the other transpositions. We also see this behaviour in measurements with IFS (not shown). [\(Zheng and Marguinaud, 2018\).](#)

Halo communication as used in semi-Lagrangian advection and global norm computation as often used in semi-implicit methods have much worse strong scalability (Figure ??). These results indicate that halo communication will become almost

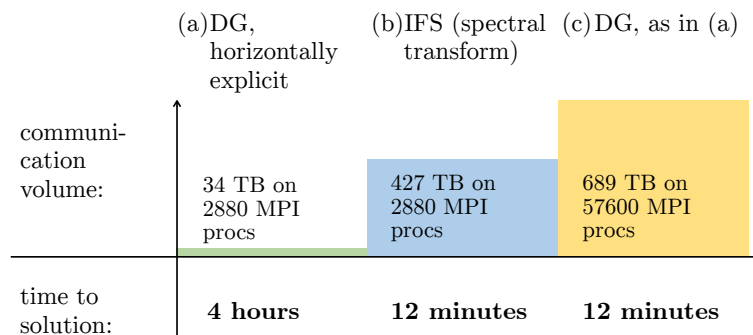


Figure 4. Overall communication volume comparing spectral element (SEM) from Müller et al. (2018) and the global spectral transform methods. The SEM requires a substantially lower amount of communication at the same number of cores, but due to the smaller timestep requires a much higher frequency of repeated communications for the given 2-day simulation. Increasing the number of MPI processes to achieve the same time to solution results in a larger amount of communication for the SEM. Here we assume SEM $\Delta t = 4s$; IFS $\Delta t = 240s$; communication volume is calculated for a 48 hour forecast SEM as 290.4kBytes per MPI task and Δt ; IFS as 216mBytes per MPI task and Δt ; $\text{IFS-SEM time-to-solution} = 20 \times \text{SEM-IFS}$ based on the performance results in Michalakes et al. (2015).

as costly as the transpositions in the spectral transform method if we use a very large number of MPI processes. An alternative which avoids transpositions and halo communication is given by the spectral element method shown in Müller et al. (2018) with explicit time integration in the horizontal direction. This leads to a very small amount of data that is communicated in each time-step because this method only communicates the values that are located along the interface between different processor domains. This method, however, requires much smaller time-steps which leads overall to an even larger communication volume (Figure 4). Figure 4 is based on the model comparison presented in Michalakes et al. (2015) and does not include all of the optimisations for the spectral element method presented in Müller et al. (2018). The spectral transform results are based on the operational version of IFS and do not contain the optimisations presented in this paper. Both models have significant potential for optimisation and it is not obvious ~~if there will be a clear winner in terms of overall communication volume~~ which method will have the lowest communication volume when fully optimised. The only true solution to avoid waiting time during communication is to overlap different parts of the model such that useful computation can be done while the data is communicated (Mozdzynski et al., 2015)(Mozdzynski et al., 2015; Dziekan et al., 2019).

3.2.1 GPU-optimisation

3.3 GPU optimisation

For the GPU version, we restructure the ~~code-prototype~~ to: allow the grid-based parallelism to be fully exposed to the GPU in a flexible manner; ensure that memory coalescing is achieved; and optimise data management. We will now describe each of these in some more detail.

The grid is a two-dimensional sphere, with a third altitude dimension represented by multiple fields at each point on the sphere. The updates are inherently parallel across this grid, so all this parallelism should be exposed to the GPU to get maximal performance. However, the original implementation had a sequential loop over one of the spherical dimensions (at a high level in the call tree of the application). We re-structured the code such that, for each operation, the loops over the three dimensions became tightly nested, and when mapping these to the GPU via OpenACC directives we used the “collapse” clause to instruct the compiler to collapse these to a single loop, such that it can map all inherent parallelism to hardware in an efficient manner.

We re-structured the code as follows:

Code example 1

```
!ACC parallel loop collapse(3)  
loop over 1st dimension  
loop over 2nd dimension  
loop over fields  
...  
Operation
```

Similarly, for library calls it is important to ~~maximally expose parallelism through use of the provided batching~~ maximise exposure of parallelism through batching computations using provided interfaces. On the GPU we perform all of matrix multiplications in the Legendre transform (2) with a single batched call of the cuBlasDgemm library. The different matrices in (2) have different sizes because the total wavenumber goes from 0 to NT. To use the fully batched matrix multiplication we pad each matrix with zeroes up to the largest size, since the library currently does not support differing sizes within a batch. This step increases the overall number of floating point operations by almost a factor 10 but still improves the overall performance (Figure 5). We perform the FFT in equation (3) with the cuFFT library, where we batch over the ~~altitude-vertical~~ dimension but multiple calls are still needed over the spherical dimension (noting FFTs cannot be padded in a similar way to matrix multiplications). Therefore the ~~code-implementation~~ remains suboptimal here: we are still not fully exposing parallelism and there would be scope for further improvements if a FFT batching interface supporting differing sizes were to become available.

We restructured array layouts to ensure that multiple threads on the GPU can cooperate to load chunks of data from memory in a "coalesced" manner. This would allow a high percentage of available memory throughput. This is achieved when the fastest moving index in the multidimensional array corresponds to the OpenACC loop index occurring at the innermost level in the collapsed loop nest ~~described above~~ as follows:




Figure 5. Roofline plot for the spectral transform dwarf at ~~TL159-125km~~ resolution ($N_T = 159$) on the NVIDIA Tesla P100 GPU. The full time-step of the original ~~code-prototype~~ is represented by the solid ~~red~~ triangle. The corresponding time-step for the optimised ~~code-prototype~~ is represented by the ~~solid-red~~ circle. Also included are partial results for kernels only (open ~~green~~ diamond) and matrix multiplication only (~~open-rectanglegreen dash~~). Each point is positioned in the plot according to its operational intensity: points under the sloping region of the roofline are limited by available memory bandwidth, and points under the horizontal region are limited by peak computational performance.

Code example 2

```

$!ACC parallel loop collapse(3)
do k=1, ...
  do j=1, ...
    do i=1, ...
      ...
      array(i,j,k)=...
    
```

In FORTRAN, arrays are structured such that elements accessed by consecutive innermost indices (i in this example) are consecutive in memory. Since i is used as the innermost loop index in this collapsed loop nest, then memory coalescing is achieved.

Sometimes matrix transposes are necessary, but where possible these were pushed into the DGEMM library calls, which have much higher-performing implementations of transposed data accesses. There remain transpose patterns within kernels involved in transposing grid point data from column structure to latitudinal (and inverse) operations, which naturally involve transposes and are thus harder to fix through restructuring. However, we optimised these using the “tile” OpenACC clause, which instructs the compiler to stage the operation through multiple relatively small tiles which can perform the transpose operations within fast on-chip memory spaces, such that the accesses to global memory are much more regular.

Data allocation on the GPU is expensive, as is data movement between the CPU and GPU. We structured the code such that the fields stay resident on the GPU for the whole timestep loop: all allocations/frees have been moved outside the timestep loop with re-use of temporary arrays, and thus all data transfer has been minimized.

The restructured algorithm achieves an overall speedup factor of 23x compared to the initial version which also used cuBlas and cuFFT but followed the CPU version more closely. Matrix multiplication performance is higher than the overall performance (in flops) and the operational intensity is increased into the compute-bound regime. Note that matrix multiplication is associated with $O(N^3)$ computational complexity for $O(N^2)$ memory accesses. The extra padding operations lead to larger N and therefore also to increased operational intensity. More details about the single GPU optimisations can be found in Mazauric et al. (2017b).

~~Going beyond~~ Beyond a single GPU ~~to multiple~~, with multiple interconnected GPUs we see a ~~massive~~ large benefit by using the modern NVLink interconnect and the recently announced NVSwitch due to the high importance of communication for the transpositions described in section 3.2. Each GPU features multiple ports of high-bandwidth NVlink connections, each providing 50 GB/s of bi-directional bandwidth when using the Volta GPUs. For full bandwidth connectivity when using more than 4 GPUs we use the NVSwitch interconnect on the DGX-2 server. The DGX-2 server has 16 Volta V100 GPUs: each with six 50 GB/s NVLink connections into the switch with routing to any of the other GPUs in the system. This allows 300 GB/s communications between any pair of GPUs in the system, or equivalently 2.4 TB/s total throughput.

When running a single application across multiple GPUs, it is necessary to transfer data between the distinct memory spaces. Traditionally, such transfers needed to be realised via host memory and required the participation of the host CPU. Not only did this introduce additional latency, but also limited the overall bandwidth to the bandwidth offered by the PCIe bus connecting CPU and GPUs. However, modern MPI implementations are CUDA-aware. This means that pointers to GPU memory can be passed directly into the MPI calls, avoiding unnecessary transfers (both in the application and in the underlying MPI implementation). This is particularly useful when using a server that features high-bandwidth NVLink connections between GPUs, in which case CUDA-aware MPI will use these links automatically. Moving our dwarf to CUDA-aware MPI gave us a speedup of 12x (Figure 6). However, even with this optimisation the all-to-all operations remained inefficient because communication between different GPUs was not exchanged concurrently. Perfect overlap was achieved by implementing an optimised version of the all-to-all communication phase directly in CUDA using the Inter Process Communication (IPC) API. Using memory handles, rather than pointers, CUDA IPC allows to share memory spaces between multiple processes, thus allowing one GPU to directly access memory on another GPU. This allowed another speedup of about 30% (Figure 6).

In Figure 7 we demonstrate how the use of DGX-2 with NVSwitch allows significantly better scaling than the use of DGX-1 for ~~the Spherical Harmonics TCo639 test case~~ a spectral transform at 18km resolution ($N_T = 639$). Note that we tune the number of MPI tasks in use: we use the NVIDIA Multi Process Service to allow oversubscription of GPUs such that, e.g. the 8 GPU result on DGX-2 uses 16 MPI tasks across the 8 GPUs (i.e. 2 operating per GPU). This is because such oversubscription can sometimes be beneficial to spread out any load imbalance resulting from the spherical grid decomposition (see below) and hide latencies. We chose the best performing number of MPI tasks per GPU in each case.

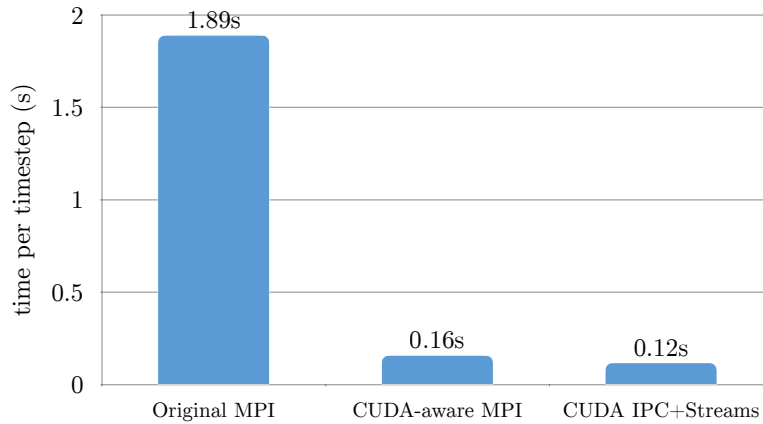


Figure 6. Computational performance of the spectral transform dwarf at ~~TCo639~~ (~~18km ensemble member~~ resolution ($N_T = 639$)) on 4 NVIDIA V100 GPUs of the DGX1 with the original MPI implementation (left), CUDA-aware MPI communication (middle) and NVLink optimised communication (right). This resolution is currently used operationally for the members of the ensemble forecast at ECMWF.

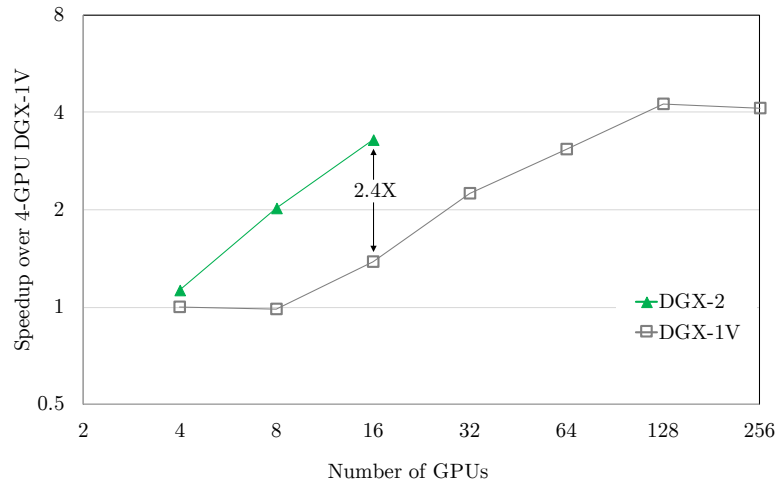


Figure 7. Computational performance of the spectral transform dwarf at 18km resolution ~~TCo639~~ (~~18km ensemble member~~ $N_T = 639$) on up to 16 GPUs on one DGX-2 and up to 32 DGX-1 servers connected with Infiniband. The DGX-1V uses MPI for ≥ 8 GPUs (due to lack of AlltoAll links), all others use CUDA IPC. DGX-2 results use pre-production hardware. The points were connected with lines for the purpose of improving readability.

As we increase the number of GPUs, the scaling on DGX-1V is limited: This is because we no longer retain full connectivity and some messages must go through the lower-bandwidth PCIe and QPI links and/or Infiniband when scaling across multiple servers. But on DGX-2 with NVSwitch, all 16 GPUs have full connectivity: that is we have maximum peak bandwidth of 300 GB/s between each pair of GPUs in use. The performance is seen to scale well out to the full 16 GPUs on DGX-2, where the difference with the 16 GPU (2-server) DGX-1V result is 2.4x. It can also be seen that the speedup going from 4 to 16 GPUs on DGX-2 is 3.2x, whereas the ideal speedup would be 4x. However, initial investigations reveal that this deviation from ideal scaling is not primarily due to communication overhead but instead to load imbalance between the MPI tasks from the spherical grid decomposition that is chosen by the application in each case, which would indicate that better scaling would be observed with a more balanced decomposition. More details about the multi-node optimisation of the spectral transform dwarf can be found in Douriez et al. (2018).

3.3.1 CPU optimisation

The First results on the supercomputer Summit of the Oak Ridge National Laboratory are shown in Figure 8. At this large scale we observe a speedup of about 2x when comparing the optimised GPU version with the initial CPU version of the spectral transform dwarf. These simulations were run at 2.5km resolution ($N_T = 3999$) and use 240 fields. In operational application the number of fields should be larger (see Section 3.1). We still need to optimise the memory consumption and the initialisation to test our prototype for more realistic numbers. Also there is still room for improving the setup of the environment on Summit, and exploring alternative programming models. The important message is that the GPU optimised dwarf from ESCAPE was capable of running on a huge number of GPUs on Summit and first results suggest that we achieve a significant speedup. We will do more analysis of these simulations in the future.

3.4 CPU optimisation

The spectral transform dwarf is based on the operational code implementation used in IFS ~~and that~~ has been continuously optimised over multiple decades. According to profiling results, it clearly appeared that the main computational intensive kernels are the FFT and matrix multiplication executed by a dedicated highly tuned library (as Intel Mathematics Kernel Library, called MKL). In support of this work we looked into different data scope analysis tools. A comparison of the different tools is available in Mazaauric et al. (2017a). The first optimisation strategy concentrated the effort on non-intrusive optimisations which have the advantage of being portable and maintainable. Among these optimisations, the use of extensions to the x86 instruction set architecture (ISA) as SSE, AVX, AVX2, AVX-512 is interesting notable, because it indicates how much of the source code can be vectorised by the compiler. When the compiler failed at vectorising some loops or loop nests, a deeper investigation of how to use compiler directives followed. As the different instruction sets are not supported by all processors, the study proposed an intra node scalability comparison study among several available systems (at the time of benchmarking).

System tuning using Turbo frequency (TUR), Transparent Huge Page (THP), memory allocator (MAP) can be done without modifying the source code. This exposes both performance gains and interesting information on dwarf behaviour. Indeed, on ATOS BullSequana x400 supercomputer with Intel® Xeon Broadwell E5-2690v4 processors, enabling turbo offers a gain

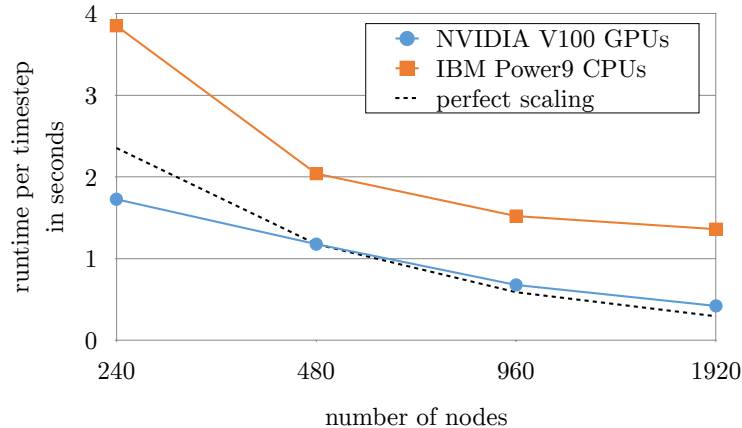


Figure 8. Strong scaling comparison between GPU optimised version and initial CPU version of the spectral transform dwarf on Summit at the Oak Ridge National Laboratory. All of these simulations are using 2.5km resolution ($N_T = 3999$) with 240 fields. The GPU version uses 6 V100 GPUs per node which leads to a maximum number of 11520 GPUs. The perfect scaling was chosen such that it goes through the GPU result at 480 nodes to illustrate the scaling efficiency of the GPU version. The points were connected with lines for the purpose of improving the readability of the plot.

equal to 11%, enabling THP gives 22%, MAP 27%, and finally the best performance (35% of performance gain) is achieved by the combination of MAP and TUR. This shows that memory management is a key point. More details about the single-node CPU optimisation can be found in Mazauric et al. (2017b).

Multi-node optimisation for CPUs focused on improving the MPI communication. The largest potential for optimisations was found to be in the preparation phase of point-to-point communications. During the preparation phase the sender side gathers the local data into a contiguous buffer (Pack operation) and hands it off to the MPI library. On the receiver side, data is then scattered from a contiguous user buffer to its correct location (Unpack operation). Pack and Unpack are nearly inevitable with scattered data because Remote Direct Memory Access (RDMA) with no gather-scatter operations are known to be often less effective, notably due to the memory pinning latency (Wu et al., 2004). It also means that sender and receiver must ~~share-exchange~~ their memory layout as they may differ. ~~In the spectral transform dwarf the Pack and Unpack algorithms were~~ The performance improvement came from reordering the loops for both pack and unpack following the memory layout of the scattered buffer. This optimisation decreased the number of tests (i.e. copy or not copy) and avoids scanning memory multiple times which was unnecessary. We reduced this with a global performance gain on the whole dwarf of about 20% (Figure 9a). A few percent of improvement came also from disabling MPI barriers which existed in the code to profile the MPI communications at low level. These barriers produced imbalance, the necessary synchronisations created contention and overall created a bias in the communications profile. The computational performance up to 30 ~~nodes on Skylake-ATOS BullSequana x400 nodes equipped with Intel® Xeon Skylake 6148 processors~~ is shown in Figure 9b. This work has been

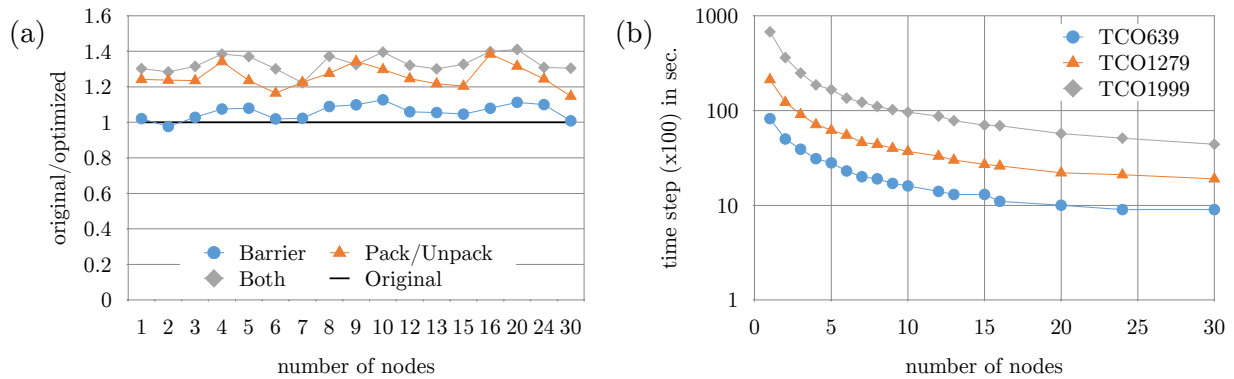


Figure 9. Performance measurements on up to 30 ~~Skylake-Xeon (SKX)-Intel@ Xeon Skylake 6148~~ nodes. Subpanel (a) shows the speed-up at ~~18km~~ resolution ~~TCo639 (18km ensemble member $N_T = 639$)~~ with regards to different optimisations of the communication preparation phase. Subpanel (b) shows the computational performance for three different resolutions of ~~TCo639 (18km ensemble member $N_T = 639$)~~, ~~TCo1279 (9km ($N_T = 1279$))~~ and ~~TCo1999 (5km ($N_T = 1999$))~~ representative of current and future operational requirements. Note the use of log-scale. ~~The points were connected with lines for the purpose of improving readability. All of these results have been obtained through measurements on the ATOS BullSequana x400 supercomputer of ATOS internal HPC resources.~~

~~performed on ATOS internal HPC resources. This optimisation can be immediately applied to the operational model IFS due to its non-intrusiveness. More details about the multi-node optimisation on CPUs can be found in Douriez et al. (2018).~~

~~Applying some of the optimisations found in restructuring the code for the GPUs to the CPU version requires some more fundamental changes which are more difficult to apply. The speedup of the GPU optimisations came from using highly~~
5 ~~optimised GPU libraries for batched DGEMM and FFT and avoiding the transposition of temporary arrays which was not necessary. On the CPU we can apply these improvements in the handling of temporary arrays as well. We used these optimisations in the fully parallelised version of the dwarf. We will continue to work on applying them to the parallel version of the dwarf. As a first step towards this goal we used them in a newly developed serial spectral transform inside Atlas which will soon be used operationally is used operationally at ECMWF for post-processing purposes since the beginning of 2019. Post-processing is~~
10 ~~run in serial mode due to the large number of concurrent post-processing jobs. Compared to the current previous operational serial transform used for post-processing we find a speedup of about 3x (Figure 10). Most of the speedup seems to be due to avoiding some of the temporary arrays and transpositions of temporary matrices.~~

3.4.1 Optical processors

3.5 Optical processors

15 ~~Optalysys have been investigating an optical implementation of the spectral transform dwarf (biFFT for limited area models as well as spherical harmonics for global models). The fundamental idea behind optical processors is to encode information into a laser beam by adjusting the magnitude and phase in each point of the beam. This information becomes the Fourier transform~~




Figure 10. Speedup of the spectral transform by porting optimisations introduced in the GPU version back to the CPUs. The base line for this comparison is the current operational post-processing library used at ECMWF. This version of the spectral transform allows the computation on limited area domains. The speedup is given here for the global transform and three examples of limited domains (Europe, UK and Hungary).

of the initial information in the focal plane of a lens. The information can be encoded into the optical beam by using spatial light modulators (SLMs) as illustrated in Figure 11. The result of the Fourier transform can be recorded by placing a camera in the focal plane of a lens. A photo of an early prototype is shown in Figure 12.

SLMs are optical devices with an array of pixels which can modulate an optical field as it propagates through (or is reflected by) the device. These pixels can modulate the phase (essentially applying a variable optical delay) or polarisation state of the light. Often they modulate a combination of the two. When combined with polarisers, this polarisation modulation can be converted into an amplitude modulation. Hence the modulation capability of a given SLM as a function of ‘grey level’ can be expressed by a complex vector, which describes an operating curve on the complex plane. Each pixel of the SLM is generally a 1-parameter device; arbitrary complex modulation is not offered by the SLM, only some sub-set. This is one of the key issues with regards to exploiting the optical Fourier transform.

Sensor arrays - essentially common camera sensors - are used to digitise the optical field. They are in general sensitive to the intensity of the light, which is the magnitude of the amplitude squared. This poses a difficulty to sensitively measuring the amplitude. Moreover, they are not sensitive to optical phase. We ~~overcome-overcame~~ this with a ~~perturbative method which determines the (constrained) optical phase from~~ method that allowed the optical phase to be derived from the intensity-only measurements. This was done by introducing known terms into the functions and measuring the resulting change in the output.

Each pixel of the SLM is addressable with an 8-bit value (256 levels). The SLM is not capable of independently modulating the magnitude and phase of the optical field. In the Optalysys processing system, the SLMs are configured to modulate both




Figure 11. Illustration of the fundamental idea behind the optical processor. The laser beam is emitted on the bottom left. Two spatial light modulators (SLM) are used together to input the complex function. The system uses beamsplitters (BS) and an optical relay to image one reflective SLM onto another, followed by a lens assembly which approximates an ideal thin lens and renders the optical Fourier transform on a camera sensor. The half-waveplate (WP) before the second SLM is used to rotate linearly polarised light onto the axis of SLM action (the direction in which the refractive index switches), thus causing it to act as a phase modulator.




Figure 12. Photo of the first prototype of the optical processor. The final product is built into an enclosure of similar size like a GPU.

the amplitude and phase in a coupled manner, such that optimal correlation performance is achieved. The optical Fourier transform and all of the functions are inherently two dimensional. The propagating light beam can be thought of as a 2D function propagating and transforming along a third direction. The system is most naturally applied to 2D datasets, and many problems can be mapped to an appropriate representation.

5 A critical aspect to realizing the potential of optical processing systems is the interface to a traditional computing platform. Bridging this gap has been a significant undertaking for Optalysys, and has resulted in the development of a custom PCIe drive board. This board interfaces to a host machine over PCIe and has direct memory access (DMA) to the system memory (RAM). It provides an interface to 4 SLMs and 2 cameras. The cameras are 4K (4096x3072). Initially, they operate at 100 Hz, but a future firmware upgrade will unlock 300 Hz operation, and 600 Hz half-frame operation, dramatically increasing the potential
10 data throughput.

There are currently two options for the SLMs. One option is using high speed binary SLMs which operate at 2.4kHz. This offers correlation at binary precision. The second option is greyscale SLMs which operate at 120 Hz. This is currently the only option to reach more than binary precision. The performance of the entire processor is determined by the part with the lowest frequency. The main bottleneck with multiple bit precision is the operating frequency of the greyscale SLM. There is currently
15 no easy solution to increase the frequency of greyscale SLMs.

Optical processing is more appropriately applied to cases where high-throughput relatively-complex operations are the priority, with less of an emphasis on numerical precision. The inherent ability of optical correlators to rapidly process convolutions naturally leads to the formation of convolution neural nets and machine learning technologies. More details about the Optalysys optical processor have been published in Macfaden et al. (2017) and Mazauric et al. (2017b).

20 **3.5.1 ~~Comparison between processors in terms of runtime and energy consumption~~**

3.6 Comparison between processors in terms of runtime and energy consumption

~~As explained in Section ?? we~~ We will lose a lot of the speedup achieved by running the spectral transform on accelerators if the CPUs are idle during the computation on the accelerators. Also we need to take the cost of data transfer between CPU and accelerator into account which has not been included in the speedup numbers in this section. To take full advantage of the
25 NVLink and NVSwitch we would need to run the entire simulation on a single node which requires at the currently operational resolutions more work on optimising the memory footprint of the model.

For the CPUs and GPUs used in this paper the overall cost is dominated by the cost of the hardware and therefore by the number of sockets/devices required to reach the desired runtime(~~see also Section ?? and the Appendix ??~~). In addition to the number of devices we also compare the energy consumption. The large number of zero operations caused in the optimised
30 GPU version by the padding of the matrices in the Legendre transform makes it impossible to do a fair comparison between CPU and GPU by comparing metrics based on floating point operations including comparing roofline plots.

In the full operational model the ~~TC639-ensemble member~~ 18km resolution ensemble member ($N_T = 639$) using 30 nodes on the Cray XC40 takes about 1.4s per time-step and the spectral transform component is about 15 percent (0.21s). Measure-

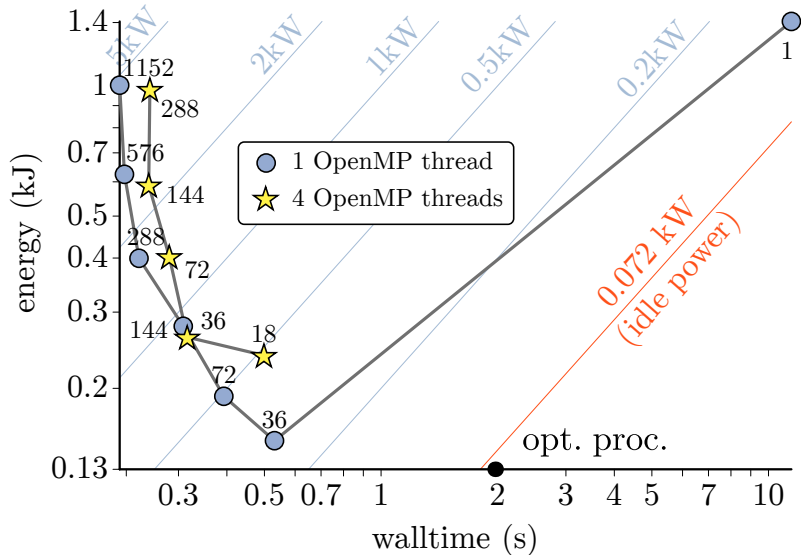


Figure 13. Log-log plot of the energy consumption vs. wall-clock time for the BiFFT dwarf and corresponding to the combination of one direct and one inverse transformation for 525 fields. Each data point is the result of averaging the outcome of two separate runs. Grey lines connect runs with the same number of OpenMP threads (1 resp. 4). Added are lines of constant power (light blue lines), including the power delivered by a node in the idle state (orange line). Indices next to each data point denote the number of MPI tasks. The black dot represents the estimate of the Optalysys optical processor when using a greyscale SLM. The performance of the optical processor at binary precision is much higher (not shown).

ments with the dwarf on the Cray XC40 at ~~resolution-TC639~~ 18km resolution ($N_T = 639$) resulted in 4.35s per timestep on a single node (4 MPI tasks, 18 threads per task), and 1.77s on 2 nodes. The energy consumption was measured at around 0.3Wh on the Cray XC40, which compares to 0.026 Wh measured on 4 V100 GPUs on a DGX1 which take 0.12s per time-step. The energy measurement on the ~~Cray-XC40~~ is based on ~~Cray-proprietary~~ power management counters. The measurement on the

5 V100 GPUs uses the nvidia-smi monitoring tool.

Tests in ESCAPE on the latest generation of Intel Skylake CPUs have shown 0.12s per timestep using 13 ~~SKX~~ Intel® Xeon Skylake 6148 nodes (connected via a fat-tree EDR network) as shown in Figure 9. This parallel CPU version has not seen the more radical changes which have been used in redesigning the algorithm for the parallel GPU and serial CPU version. There might still be potential for more substantial optimisations in the parallel CPU version which we will explore in future research.

10 A comparison between CPUs and optical processor with greyscale SLM is shown in Figure 13. The energy consumption of the optical processor is much lower than for the CPU. The runtime of the optical processor is larger due to the relatively slow performance of the grayscale SLM which is currently necessary to reach the precision necessary for NWP applications. More details about the comparison between CPU and optical processor can be found in Van Bever et al. (2018a).

3.7 Elliptic solver Sustainability of code optimisation techniques

3.7.1 Background

The dwarf originates from the elliptic solver used in the semi-implicit time integration of IFS-FVM (Smolarkiewicz et al., 2016; Kühnlein et al., 2016).

We employ the Generalised Conjugate Residual (GCR, Eisenstat et al., 1983) approach to solve the following linear elliptic

5 problem

$$\mathcal{L}(\psi) = \sum_{I=1}^3 \frac{\partial}{\partial x^I} \left(\sum_{J=1}^3 C^{IJ} \frac{\partial \psi}{\partial x^J} + D^I \psi \right) - A\psi = Q,$$

with variable coefficients A, C^{IJ}, D^I and rhs Q , assuming either periodic (Szmelter and Smolarkiewicz, 2010) or Neumann boundary conditions. This dwarf intends to explore preconditioning strategies, where (??) is augmented to $\mathcal{P}^{-1}[\mathcal{L}(\psi) - Q] = 0$,

with \mathcal{P} a linear preconditioning operator that approximates \mathcal{L} but is easier to invert. Given a suitable preconditioner, this

10 auxiliary problem can converge faster due to a closer clustering of the eigenvalues with the superposition of \mathcal{P}^{-1} and \mathcal{L} (Thomas et al., 2003). More details about the elliptic solver dwarf can be found in Mengaldo (2016).

3.7.1 Multigrid Preconditioner

A challenge for the elliptic solver is to find an effective preconditioner \mathcal{P} which is a good approximation to the linear operator \mathcal{L} and is more economical to solve. The inversion of \mathcal{P} can be simplified by two adaptations of \mathcal{L} ; firstly by taking the matrix C in

15 (??) to be diagonal and secondly by separating the inversion of the unstructured horizontal and structured vertical components of \mathcal{P} —mirroring the natural anisotropy of the terrestrial atmosphere. Solutions to \mathcal{P} are sought using iterative procedures akin to those in appendix B of Thomas et al. (2003). Such solvers are quick to damp errors associated with length scales of the underlying discrete meshes on which the computations are performed, but can be slower damping errors associated with longer length scales of the horizontal domain. The basic idea behind the multigrid preconditioner is to use a nested tower of grids

20 of varying resolution to quickly eliminate solution errors associated with those grids. Each subsequent coarser grid evaluation iteratively solves the residual problem of the previous finer grid level, which provides an error correction to the finer grid solution in a V-cycle configuration (Figure ??).

To solve the problem at a coarser resolution the field needs to be restricted to the coarser mesh. Restriction and interpolation arrangements on the octahedral mesh for the multigrid elliptic solver dwarf. Each coarsening step incorporates a smoothing

25 step, which minimises the solution errors associated with the given resolution. The coarsest grid correction is usually found using some sort of direct solve technique which may be too costly to perform at finer resolutions. This solver may or may not match the smoothing method employed on the intermediate grids. After obtaining an estimate on the coarsest mesh, the errors are interpolated back to the finest mesh via each intermediate mesh. The coarser grid error is used to correct the solution error on a finer grid at each stage. This may require an additional smoothing step to be performed if the coarser grid correction

30 reintroduces errors associated with the finer resolution. The whole cycle might employ some form of smoother/solver to both initialise and finalise the solution. The preconditioner utilises either a vertically implicit Richardson or weighted line Jacobi

method as both smoother and solver for the horizontal inversion of the preconditioner operator. The vertical operator is inverted using a bespoke tridiagonal solver, since the vertical component has a structured data arrangement which is always local to the same process.

5 The employed mesh coarsening strategies utilise the Atlas framework for mesh generation and create a tower of nested octahedral meshes, where all nodes on the coarser grids correspond to nodes on the finer grids. The nested mesh structure maintains the parallel decomposition.

First results for the baroclinic instability benchmark show a significant reduction in the number of iterations of the elliptic solver (Table ??). Work on optimising the multigrid preconditioner is still in progress. Performance measurements will therefore follow in a later publication. More information about the multigrid preconditioner can be found in Müller et al. (2017).

10 grid-O180-O360-O720 resolution-61km-31km-16km iterations-without-multigrid-11-15-26 iterations-with-multigrid-6-7-11
Number of iterations at three different resolutions for the baroclinic instability test case.

Furthermore, the unstructured-mesh weighted line Jacobi method developed as a smoother for multigrid can be employed as a single-level preconditioner, with a significantly improved convergence of GCR compared to the Richardson preconditioner; see also Kühnlein et al. (2018) for discussion.

15 3.8 HEVI time-integration

As part of the ESCAPE project a review paper of different time discretisation strategies for NWP applications has been published in Mengaldo et al. (2018) since time integration affects cost considerably. In this review we found that for hydrostatic models like the currently operational IFS using semi-Lagrangian advection with semi-implicit time integration is extremely well suited and very difficult to beat due to its superior performance. At some point in the future we expect to reach resolutions
20 at which we will need to use Eulerian-based time integration (EBTI), like for example horizontally explicit vertically implicit (HEVI) schemes. As an example of this class of time integration methods we worked on Runge-Kutta implicit-explicit (RK-IMEX) methods. The class of RK-IMEX time discretization schemes may be defined as follows. First a partitioning of the right hand side (RHS) terms of the system to be solved is introduced through

$$\partial_t X = \mathcal{E} + \mathcal{I},$$

25 where the term \mathcal{E} denotes the part of the system RHS to be treated explicitly, and \mathcal{I} the part of the system RHS to be treated implicitly. Then, two different multi-stage RK schemes are respectively applied to \mathcal{E} and \mathcal{I} parts. The RK scheme applied to \mathcal{E} is purely explicit whereas that applied to \mathcal{I} allows implicit evaluations at each sub-stage. The result may be written under the general form

$$\frac{X^{(j)} - X^0}{\Delta t} = \sum_{i=1}^{j-1} \tilde{a}_{ji} \mathcal{E} [X^{(i)}] + \sum_{i=1}^j a_{ji} \mathcal{I} [X^{(i)}],$$

$$\frac{X^+ - X^0}{\Delta t} = \sum_{j=1}^{\nu} \tilde{b}_j \mathcal{E} [X^{(j)}] + \sum_{j=1}^{\nu} b_j \mathcal{I} [X^{(j)}],$$

where $\nu \geq 2$ is the total number of sub-stages of the RK-IMEX scheme, i, j are integer indexes such as $1 \leq i \leq j \leq \nu$, $X^{(j)}$ denotes the value of the state variable at the j -th sub-stage, and the superscripts "0" and "+" correspond to the values of the state variable at times t and $t + \Delta t$, respectively. Notations like $\mathcal{E} [X^{(i)}]$ indicate that the terms of the sub-system \mathcal{E} are evaluated using the state variable at sub-stage $X^{(i)}$.

The coefficients $\mathcal{A} = (a_{ji})$, $\tilde{\mathcal{A}} = (\tilde{a}_{ji})$ for $(i, j) \in [1, \nu] \times [1, \nu]$, and the weight-vectors $(b_j, c_j = \sum_{i=1}^j a_{ji})$ and $(\tilde{b}_j, \tilde{c}_j = \sum_{i=1}^j \tilde{a}_{ji})$ for $j \in [1, \nu]$ may be classically represented by a double Butcher tableau: $\tilde{c} \tilde{\mathcal{A}} \tilde{b}^T \quad \quad \quad c \mathcal{A} b^T$

The first Butcher tableau defined by $(\tilde{\mathcal{A}}, \tilde{b}, \tilde{c})$ describes the explicit part so that $\tilde{a}_{ij} = 0$ [The optimisation results presented in this section are just an example](#) for $i \geq j$, and the second one (\mathcal{A}, b, c) corresponds to the implicit part of the scheme. RK-IMEX schemes, using such a double Butcher tableau, are traditionally labelled in literature with the nomenclature $[\text{NAME}]_k(s, \sigma, p)$, where k denotes the order of accuracy of the explicit part, s , the number of implicit inversions to be performed in the implicit part (i.e. the number of non-zero diagonal coefficients in \mathcal{A}), σ , the storage factor (i.e. the minimal number of explicit sub-stages that need to be stored to complete the time-step), and p , the overall order of accuracy of the scheme. The particular RK-IMEX schemes that have been analysed in this work, are the so-called UJ3(1,3,2), ARK2(2,3,2), and Trap2(2,3,2) scheme for the finite-difference (FD) or finite-volume (FV) discretisation and additionally SSP3(3,3,2) and SSP3(4,3,3) for a Discontinuous Galerkin (DG) method, according to this nomenclature. The double Butcher tableau for each of these schemes is given in Lock et al. (2014).

The most suitable RK-IMEX HEVI scheme for time-critical NWP applications is the one that achieves the best balance between low computational cost and good stability property. In this prospect, the RK-IMEX time-schemes UJ3, ARK2, and TRAP2, identified in the literature as promising time-discretisation in HEVI context, have been analysed in more detail taking into account the stability effects of the advection processes and orographic forcing terms introduced by the use of terrain-following coordinate. Drawing from all these analyses, it firstly appears that RK-IMEX HEVI approaches based on only two Butcher tableaux might not be optimal for dealing with the different dynamical processes involved in a fully compressible system, and their multiple interactions, and notably the feedback between the horizontal divergence and pressure gradient terms. Secondly, in presence of steep sloped orography, the use of the covariant horizontal wind component as a prognostic variable has been found to be beneficial for the stability of all examined RK-IMEX HEVI schemes. Exploiting all these results, a optimal HEVI scheme, termed as TRAP2 "covariant-implicit", has been designed. This newly proposed scheme offers the largest stability limit while being as cheap as the original TRAP2 scheme in terms of computational cost. As a proof-of-concept, the suitability of this scheme has been confirmed experimentally in a complete fully-compressible vertical plane model. Figure ?? shows the improvement of the stability region when compared with the traditional scheme. More details of this work have been published in Colavolpe et al. (2017). Stability region of (a) common Trap2 time integration scheme and (b) enhanced

TRAP2 time integration scheme (Colavolpe et al., 2017). The vertical axis of the stability graphs show the time step while the horizontal axis represents the advection velocity. Values larger than 1 mean that the scheme is unstable. As we can see the new scheme provides a significantly improved stability region. A HEVI formulation of the general DG approach has been implemented, too. Problems occurring with the splitting of the numerical flux and the formulation of proper boundary conditions have been solved for simple test problems as the linear advection and the linear wave equation. For the more upwind-biased numerical fluxes used in DG, the use of the SSP3 schemes seems to be more appropriate.

3.8 Semi-Lagrangian advection

3.7.1 Background

We consider a three-dimensional semi-Lagrangian advection scheme. Advantages of semi-Lagrangian schemes are stable and accurate integration with long time steps and efficiency for multiple tracers. The semi-Lagrangian transport scheme solves the following evolution equation

$$\frac{D\psi}{Dt} = 0 \quad \text{with} \quad \frac{D}{Dt} = \frac{\partial}{\partial t} + \mathbf{v} \cdot \nabla,$$

where ψ is a scalar field advected by the wind $\mathbf{v} = (u, v, w)$.

To solve Eq. we integrate along the trajectory of a fluid parcel in the time interval $[t, t + \Delta t]$ which yields $\psi_a^{t+\Delta t} = \psi_d^t$. The subscripts a and d denote the arrival and departure points of the flow trajectory, respectively. The arrival point is the location of a parcel at time $t + \Delta t$ and coincides with a grid point. The departure point represents the parcel's location at time t , and typically lies somewhere in the space between grid points. Hence, the departure point has to be found. The semi-Lagrangian scheme solves for each arrival point with coordinates \mathbf{r} the trajectory equation

$$\frac{D\mathbf{r}}{Dt} = \mathbf{v}$$

to determine the departure points \mathbf{r}_d , and thereafter remaps the field ψ to the set of these points. In the ESCAPE project, we implemented a prototype for the semi-Lagrangian advection using Atlas. To solve the trajectory equation an iterative method based on the second-order mid-point integration scheme is used (Hortal, 1999; Temperton et al., 2001). For the interpolation, the LAITRI (LAgrangian Interpolation TRilinear) procedure of IFS is adopted. More details about these two dwarfs can be found in Mengaldo (2016).

3.7.1 Optimisation for CPUs and Xeon Phi

Work on optimising semi-Lagrangian advection focused on optimising the LAITRI dwarf. LAITRI is a heavily-used subroutine in the European IFS global weather modelling system. It accounts for about 4% of the total runtime of IFS. The work on optimising this dwarf was done at the beginning of the ESCAPE project when Knights Landing (KNL) was new, and the interesting question was how KNL would compare with the older Ivy Bridge CPU and Knights Corner (KNC) accelerators.

Strong-scaling studies of the LAITRI dwarf over increasing OpenMP thread count have shown that best performance is achieved on the Ivy Bridge node by using 40 OpenMP threads and 60 threads on KNC and KNL. We compare the best time-to-solution for different hardware settings, and with thread count being maintained as constant across each platform, at the optimum point of the scaling plots. We experimented with the following settings: without automatic vectorisation (by 5 compiling with the `-no-vec` switch in ifort), with vectorisation, and with data alignment at 64 byte boundaries. The runtime measurements are shown in Figure ???. Time to solution of LAITRI dwarf for Ivy Bridge, Knights Corner and Knights Landing processors. The results show that improvements in time-to-solution are incremental for any given optimization on a fixed platform, with no particular setting giving considerable speedup compared to others. However, there is a marked performance boost on the KNL platform over both KNC and the Intel Ivy Bridge Xeon processors tested. This is a turnaround from previous 10 work comparing KNC to Ivy Bridge, where it was found (with similar results elsewhere in HPC) that the KNC performance was disappointing. Not only does KNL surpass KNC (by an order of magnitude in our tests), it also outperforms Ivy Bridge by 2 to 1 in time to solution. The settings to achieve this are straightforward, involving 3 main ideas – correct use of OpenMP (e.g. ensuring correct ‘first touch’ of data), suitable compiler data alignment directives to facilitate good vectorisation performance, and prudent setting of runtime variables such as `KMP_AFFINITY` and `KMP_HW_SUBSET`. More details about this work can 15 be found in Robinson et al. (2016).

3.8 MPDATA advection

3.7.1 Background

MPDATA refers to a class of non-oscillatory forward-in-time high-resolution numerical schemes for the advective terms in flux-form formulations of fluid equations. In contrast to the semi-Lagrangian approach based on Eq. , MPDATA solves the 20 advective transport problem in the form of an Eulerian conservation law

$$\frac{\partial \rho \psi}{\partial t} + \nabla \cdot (\mathbf{v} \rho \psi) = 0 ,$$

where ρ represents the fluid density. MPDATA schemes are at the basis of the Finite-Volume Module of IFS (henceforth IFS-FVM), a novel dynamical core formulation under development at ECMWF (Smolarkiewicz et al., 2016; Kühnlein et al., 2018). The finite-volume MPDATA implementation in IFS-FVM is described in Kühnlein and Smolarkiewicz (2017), including a 25 comprehensive list of references on MPDATA methods and their broad applicability.

The basic principle of MPDATA is most suitably described as an iterated upwind (alias donor cell) scheme: The initial iteration represents the first-order accurate upwind scheme with the advective velocity given by the physical flow. Subsequent iterations are also based on the upwind scheme, but the updated field is advected with a properly defined pseudo-velocity designed to compensate to selected order (typically second) the spatial and temporal truncation errors of the previous iteration. 30 The resulting scheme is at least second-order accurate in time and space, fully multidimensional and conservative. Due to the consistent application of the upwind differencing, MPDATA retains the characteristic features of a relatively small phase error and strict sign preservation of the transported field. Various extensions of the basic MPDATA scheme are available, such

as for the incorporation of arbitrary right-hand-sides, the transport of fields with variable sign, and the nonoscillatory option that ensures solution monotonicity. Moreover, structured-grid flux-form finite-difference and unstructured-mesh finite-volume formulations of MPDATA exist, see Smolarkiewicz and Margolin (1998); Smolarkiewicz and Szmelter (2005); Kühnlein and Smolarkiewicz (2017). More details about the dwarf can be found in Müller et al. (2017).

5 3.7.1 CPU optimisation

Profiling with a solid-body rotation on a $O1024$ mesh indicates that for a large number of computing nodes most of the runtime is spent in MPI communication. For this reason we focused our CPU optimisation on optimising MPI communication.

The MPI library, especially the Intel MPI one, relies on several interconnection protocols. It has to be wisely chosen according to the cluster hardware. In the benchmark cluster the interconnect network is based on the Infiniband protocol.

10 Traditional Infiniband support uses the Reliable Connection (RC) protocol to exchange MPI messages, but the User Datagram (UD) protocol has emerged as a lower memory consumption, more scalable alternative. The first optimization has been to enable this protocol in the Intel MPI library.

A second optimization is the replacement of the “manual” implementation of the AlltoAllV algorithm inside the `halo_exchange` subroutine by the `MPI_AlltoAllV` function. To enhance the scalability, one may consider overlapping the MPI communications with computation loops. This has been introduced by two variants based on a single core idea: the computation loops have to be evaluated on the data to be exchanged before sending them and then they are evaluated on the private data while the communications are received. The first one is implemented via a mask array separating the shared versus the private data. The second one is implemented by an indirection array storing the shared (resp. the private) indices of the data.

The modified versions have been launched on a cluster with nodes based on Skylake processors (16 cores, 192GB) interconnected by an Infiniband EDR network. Figure ?? shows scalability results of multi-node optimizations on a 16 physical cores SKX with 1 MPI task per socket and 1 OpenMP thread per physical core. The two combined optimizations (UD Infiniband protocol and AlltoAllV) show a speedup of nearly 15% (20s vs 23s on 64cores, 12s vs 14s on 128 cores). The two variants of the asyene patterns are less performant than the AlltoAllV implementation. In the first variant, the loops are executed two times with a conditional inside them. This alters the performance. The second option leads to non-contiguous data which is also less efficient. MPDATA scalability results of multi-node optimisations on a 16 physical cores SKX with 1 MPI task per socket and 1 OpenMP thread per physical core. The two combined optimisations (UD Infiniband protocol and AlltoAllV) show a speedup of nearly 15% (20s vs 23s on 64cores, 12s vs 14s on 128 cores). The two variants of the asyene patterns are less performant than the AlltoAllV implementation.

3.7.1 GPU optimisation

30 Previous work on optimising MPDATA for GPUs considered the finite-difference formulation on structured meshes (Wyrzykowski et al., 2017). In the ESCAPE project, we focused on the finite-volume formulation of MPDATA which supports unstructured meshes. As a test case we use a solid-body rotation over the pole on an octahedral mesh $O128$ with three levels in the vertical. The most computationally expensive kernel beside communication was identified as “`compute_fluxzdiv`” which computes the divergence.

This was found through profiling on the CPU where it is responsible for 25% of overall runtime. There exist three nested loops in this kernel which must be mapped to the parallelism of the GPU. The extents of these, for the test case in use, are: L1 loop over 71424 nodes, L2 loop over 3 levels and L3 loop over 213199 edges. The resulting performance of the initial implementation on the GPU was low: it is only able to achieve 44GB/s data throughput on the NVIDIA Tesla P100 GPU, which is less than 10% of that achieved by the STREAM benchmark. The reasons for this are related to suboptimal parallel decomposition and data layout, and the accessing of data in deep structures.

There are a number of possibilities for how the parallel mapping can be implemented and the original choice was suboptimal. The reason for this is that the loop assigned to CUDA threads within each block (at L2) has an extremely small extent, where typically we need much more parallelism at the CUDA thread level to make good use of the vector nature of the CUDA execution model. Instead, we choose to collapse the two outermost loops (through use of the OpenACC parallel loop collapse(2) directive) and assign the parallelism across this collapsed loop to both CUDA blocks and threads within each block. This allows the compiler to decide a much more suitable extent of vectorisation. An OpenACC loop seq directive is applied to the innermost loop, such that each thread will perform all of this loop in a sequential manner (satisfying the requirements of the reduction). With this new parallelisation strategy, it is important to ensure that data is accessed in a coalesced manner for the different arrays, in order to achieve a high percentage of memory bandwidth. For coalescing, we need consecutive threads (corresponding to consecutive indices of the vertical level) to access consecutive memory addresses, and fortunately the original data layout of these arrays (with the level being the fastest moving innermost index in Fortran) already satisfies this requirement, so no further data layout modifications are necessary.

The kernel accesses several read-only data elements and structures. For these, best performance is achieved when the compiler maps the data to the fast on-chip constant cache on the GPU. However, we find that, for the case of the deep array access like this %geom%node2edge_sign the compiler does not make full use of this capability. But, if we copy this to a regular “flat” array, ahead of kernel execution and use this in place of the original structure we see an increase in constant cache utilisation and improved performance. Furthermore, we can see that the operation involves division by a constant. We replace this by multiplication by the reciprocal of the constant (calculated in advance), which further boosts performance.

These optimisations decrease the time taken by the kernel by a factor of 9.4x. The achieved throughput of the optimised version is measured by the NVIDIA profiler to be 344GB/s, which is 66% of the value measured using STREAM benchmark, indicating that we are reasonably close to the hardware limit. More information about the optimisation of this kernel can be found in Mazaauric et al. (2017b).

Optimising all of the MPDATA kernels in a similar way gives us an overall speedup of 57x compared to the initial OpenACC port of the CPU version which brings the dwarf close to the roofline (Figure ??). This includes optimised data management such that the fields stay resident on the GPU for the whole timestep loop: all allocations/frees have been moved outside the timestep loop with temporary work arrays being re-used, and all host/device data transfer has been minimized. Roofline plot of the original and optimised versions of MPDATA for the solid body rotation over the pole at O512 resolution on the NVIDIA Tesla V100 GPU. The points are positioned in the plot according to their arithmetic intensity: points under the sloping region of the

roofline are limited by available memory bandwidth, and points under the horizontal region are limited by peak computational performance. These results were obtained with PGI 17.9, CUDA 8.0.61 and OpenMPI 1.10.7.

When utilising multiple GPUs, two halo exchanges are required each timestep for each of two fields. The subroutine for the halo exchange is provided by the Atlas library, which expects the data to be resident on the host CPU. Therefore, the OpenACC directives are required to copy the array from GPU to CPU before the operation, and back to the GPU after. As will be seen, this data movement has a huge overhead. Therefore, we developed a new custom halo exchange mechanism to be used in place of the original, which allows data to be kept resident on the GPU throughout, which we will now describe.

With structured grids, halo exchanges are relatively straightforward since the halo data on each subdomain corresponds to a small subset of “neighboring” subdomains in a clear manner. However, the data structures in this dwarf are unstructured grids, and as a result each subdomain may have halo data elements corresponding to any of the other subdomains. Therefore, the halo exchange requires an “all-to-all” communication pattern, where the sizes vary. We implement this as follows:

- Pack “send” buffer containing “edge” data on each GPU (using OpenACC), for each corresponding remote GPU.
- Exchange this data using CUDA-aware MPI_AlltoAllV (like in the GPU optimised spectral transform dwarf).
- Unpack “receive” buffer on each GPU (using OpenACC), with each element being stored in its appropriate halo location.

Figure ?? shows performance measurements with this approach. The dependence of the MPDATA time-step time, for for the solid body rotation over the pole at O1024 resolution, on the number of V100 GPUs in use within a single DGX-1V server. Orange circles denote the original halo exchange communication mechanism which involves host-device data transfer and the host-based Atlas library. Green triangles denote use of a new custom CUDA-aware halo exchange mechanism. When only a single GPU is in use, no halo exchange is required, and the runtime per time-step is measured (using the optimisations described above) to be 0.0214s. But when we enable the halo exchange (still on a single GPU), using the pre-existing communication mechanism, then host/device data transfers occur for the entire field which has a huge overhead, and the time increases by an unworkable factor of 3.9x to 0.0835s. With our new mechanism, the overhead is only around 1% with the time at 0.0216s. By keeping data resident on the GPU we have removed the overhead and allowed effective multi-GPU scaling.

3.7.1 Domain Specific Language

The GridTools framework provides a set of tools for developing numerical methods of weather and climate applications. The emergence of new computing architectures and accelerators in the supercomputing systems where weather and climate applications are run pose a challenge to efficiently maintain and run weather models. Typically weather models are complex systems with large codebases (from hundred thousands to millions lines of code). Differences in the computing architectures make adapting models to new architectures a daunting task. Often, different architectures offer different memory spaces that must be managed explicitly (like GPUs), efficient computations on gridded fields require storing the multidimensional fields with different memory layouts, etc. And the nested loops over dimensions and performance optimisations (such as tiling/loop

blocking, loop fusion, etc.) are specific to each hardware architecture. Additionally, they might require the use of different programming models.

one of the dwarfs. Similar improvements have been achieved for the advection scheme MPDATA (Douriez et al., 2018) and the radiation scheme ACRANEB2 (Poulsen and Berg, 2017). These optimisations are hardware specific and will be difficult to maintain on upcoming new architectures. As a strategy towards performance portability and sustainability we have worked in the ESCAPE project on domain specific languages through the GridTools framework. The main goal of the GridTools library is to provide a solution for weather and climate models to run one code base on many different architectures (portability) and achieve good performance (performance portability). However, the main operational product of GridTools so far focused on solutions for lat-lon grid models like COSMO. The work developed in the ESCAPE project aimed at extending the DSL support for irregular grids and the efficient generation of backends for multiple architectures.

In order to index fields in their corresponding location and to establish the connectivity with neighbors, the GridTools backend for irregular grids introduces the concept of a location type, which can be edges, cell centers and vertices for any type of elementary shape. The location type can be used in the declaration of the storages that hold fields as well as in the operators.

Another concept of the GridTools backend for irregular grids is the topology of the grids, which knows the connectivity among the grid points in each location type for each of the supported grids.

Some irregular grids, like the octahedral/icosahehedral grid will still retain a structure which allows deriving rules for extracting the connectivity of the topology without the use of unstructured meshes. Other grids, like a generic reduced Gaussian have a structure, although more complex, and do not allow to derive easy methods to extract the connectivity. The latter group will exhibit a totally irregular pattern and will require the use of an unstructured mesh.

Memory data layouts of the fields where computations operate are crucial for performance, particularly on modern accelerators. Modern CPU processors operate on large vector widths (like AVX-512) while accelerators like NVIDIA GPUs compute on warps of 32 CUDA threads. Both provide more efficient use of memory if loads and stores are performed on aligned and coalescing accesses. A coalescing access will require that the memory loads/stores of different parallel cores of a GPU warp or elements of a vector instruction in a modern CPU processor are contiguous in memory.

This is straightforward to achieve for lat-lon grids, with memory layouts organized in rows/columns for the (i,j) indexing space. On the contrary, irregular grids without any regular pattern in the grid do not allow to obtain coalescing accesses. However, many grids employed by weather and climate models are derived from platonic solids and retain their original structure. Examples are the octahedral grid (widely used in the ESCAPE dwarfs), the icosahehedral grid or the cubed sphere.

In the ESCAPE project, we explored how to obtain structured domain decompositions for the octahedral grids and conducted performance comparisons of different types of indexing on accelerators, with the goal to determine basic properties and optimal memory layouts for optimal performance of the ESCAPE dwarfs. Figure ??b shows a possible domain decomposition of the octahedral grid that preserves the original structure of the octahedron. (a) Equal partitioning domain decomposition of the octahedral grid. (b) Structured domain decomposition based on parallelograms (as seen from two different angles):

With this domain decomposition, a new indexing method that uses colouring for cells/edges and vertices was introduced, with the important property that all the loads/stores of the computational patterns found in the dwarfs are coalescing and a large fraction of them are also aligned (Figure ??). Structured indexing of (a) cells, (b) vertices and (c) edges of a parallelogram of the icosahedral/octahedral grid. Gray cells indicate padding inserted in order to align accesses to vertices within the compute domain. Number of colors is 1 for vertices, 2 for cells (downward / upward triangles) and 3 for edges. Each cell is indexed with a tuple (row, color, column).

Figure ?? shows that the structured numbering yields best performance when combined with the indirect addressing. The DSL backend of GridTools for irregular grids developed within the ESCAPE project implements structured direct addressing and irregular indirect addressing. Future developments will support structured indirect addressing which gives best performance on GPU accelerators for grids that contain a structure and irregular indirect addressing using a Hilbert space filling curve that provides still a good bandwidth for those grids for which a structure cannot be exploited (e.g. following coastlines). Bandwidth (required data transfer / computation runtime) in GB/s of two stencil computations for the different indexing and access methods on an octahedral patch of size 128x128x80 for a P100 GPU. These results show that the structured numbering yields best performance when combined with the indirect addressing. The DSL backend of GridTools for irregular grids developed within the ESCAPE project implements structured, direct and unstructured, indirect addressing. Future developments might support structured, indirect addressing which gives best performance on GPU accelerators for grids that contain a structure and Hilbert curve, indirect addressing that provides still a good bandwidth for those grids for which a structure cannot be exploited.

These DSL developments have been used to implement a portable version of the MPDATA dwarf. The DSL version hides such details as the nested loops and the OpenACC directives used to specify properties of the GPU kernel and data layouts of the FORTRAN arrays. Furthermore, the DSL allows to compose several of these operators together, which is used by the library to apply advanced performance optimisations like loop fusion or software managed caches.

Since most of the weather and climate applications are memory bandwidth bound on modern processors and accelerators, many of the performance optimisations focus on the best utilisation of the memory subsystem of the computing architectures. In order to optimise memory bound kernels, one of the most prominent optimisations is the combination of tiling and loop fusion that increases data locality. All computing architectures offer a memory system with different levels of cache or scratch pad. Since the bandwidth of a cache level is typically orders of magnitude larger than main memory, the use of the cache of the memory system to reduce main memory accesses increases significantly the performance of the memory bound applications. Architectures like traditional CPUs or Intel XeonPhi have an automatic caching mechanism that does not require explicit instructions at the software level. However techniques like tiling or loop fusion are crucial in order to fit temporary computations into the fastest levels of the cache system.

On the other hand, NVIDIA GPUs require an explicit declaration in the programming model for the use of the different levels of the cache system, like the shared memory.

The composition of stages of the DSL allows the library to apply these loop tiling and fusion. In the MPDATA example shown for the computation of the fluxes, multiple fields are reused between the different computations. The theoretical calculations give us a number of 1140638 main memory accesses without fusion and a number of 357120 memory accesses with fusion.

Measurements on an Intel Haswell E5-2690 CPU demonstrate that the time per grid point update of the computation of the fluxes of the MPDATA dwarf is reduced by almost 50%. Among other optimisations, the use of a DSL allows to fuse all the stages that form a single computation of the MPDATA, using high bandwidth scratch pad for intermediate variable, which increases the data locality of the algorithm. Such optimisations can only be performed since the library assumes a parallel model that supports only specific and limited computational patterns that can be expressed by the DSL, as opposed to general purpose language compilers that cannot make such assumptions. Comparing the Fortran [Comparing the FORTRAN OpenACC kernel with the DSL version](#) gives us a speedup of 2.1x for the DSL version. This speedup could also be achieved by hand-tuned optimisation. The DSL prevents the repeated manual effort of tuning the code for multiple architectures. At the same time the DSL allows to perform optimisations which would otherwise make the code unreadable. More details about this work including code examples on how to use the new backend to GridTools can be found in Osuna (2018). [More work on sustainability through domain specific languages will follow through the ESCAPE-2 project which started in October 2018.](#)

3.8 High-order finite difference

As an alternative to the spectral transform and finite volume discretisation we created a global shallow-water model named GRASS (Global Reduced A-grid Spherical-coordinate System), which uses high-order finite differences while still supporting the reduced octahedral grid that is used operationally in IFS and ALARO. This means that the number of grid points decreases with increasing latitude.

Longitudinal derivatives are computed using Lagrange Sine/Cosine representation with a one-dimensional stencil along the longitudinal row. A Fourier representation along longitudes is also possible near poles. Meridional derivatives are calculated by first interpolating values at remote latitudes along longitudes in order to get a meridionally-aligned set of data at each grid-point, as illustrated in Figure ??, and then a one-dimensional discrete A stencil used to calculate a fourth order meridional derivative. meridional derivative operator is applied. The longitudinal interpolations are made with Lagrange Sine/Cosine representation, which takes advantage of the longitudinal periodicity to give a more accurate interpolation if the stencil is spanning a large part of the circle. The meridional derivative is calculated by applying the classical centred derivative to the interpolated values.

This approach was implemented and parallelised with MPI and OpenMP. Solid-body rotation experiments were carried out to make sure there were no discretisation errors for this flow, and to show that the scheme can be stable. Drawing from these results most favourable configurations were identified.

GRASS is specifically designed to achieve the best quality for complex and challenging flows at high resolutions on the sphere. Since the exact solutions are not known for these flows, the references are taken from simulated solutions recognised as of the best quality in the scientific literature. For this purpose the most challenging case to date for a two-dimensional model on the sphere was chosen. The reference is the converged solution presented in Scott et al. (2015). Our results demonstrate very good agreement with the reference solution which uses a spectral transform method (Figure ??). More details about this work will be in a pair of papers submitted for publication (Bénard and Glinton, 2019; ?). High order Finite Difference result for the shallow water test case from Scott et al. (2015) at day 6 and resolution N5761 (i.e. 5761 latitudinal rows with approximately 3.5 km horizontal grid spacing):

4 Physics dwarfs

Physics parameterisations account for about 30% to 40% of the overall forecast runtime (compare Figure 3). Out of the large number of different parameterisations we focused on the two computationally most expensive parameterisations: a radiation scheme called ACRANEB2 from ALARO and a cloud microphysics routine called CLOUDSC from IFS. Both of these two dwarfs have many features of physics modules in weather and climate models in general. Thus, they include frequent usage of transcendental functions, and complex loop and conditional structures. This makes them interesting also in a broader context.

3.1 Radiation scheme ACRANEB2

The radiation schemes in numerical weather prediction and climate models take up a considerable amount of the overall running time of these models. Here, “radiation” is implicitly taken to mean *electromagnetic radiation*. The heating due to absorption of shortwave (solar) and longwave (terrestrial heat) radiation is the initial driver of all atmospheric processes with the exception of volcanic events.

Neither the shortwave nor the longwave radiative transfer can be solved within a reasonable amount of time from basic principles. In addition to the spatial and temporal approximations that are necessary to make for all physical processes in atmospheric models, it is necessary to make approximations in the spectral dimension and the directional dimensions. Thus, it is not feasible to calculate the radiative transfer for each absorbing and emitting line of the atmospheric gases; instead a limited number of spectral bands are defined for which the radiative transfer is calculated. For shortwave irradiance the radiative transfer in most current models is only considered for the direct solar beam, upward diffuse irradiance and downward diffuse irradiance. Thus, the complex directional variability of shortwave irradiance is not considered. This is called the two-stream approximation. For the longwave irradiance the two-stream approximation is also used in most current models. To sum up, many approximations are currently made in order to calculate radiative transfer in weather and climate models, and even with these approximations they are still computationally very expensive.

Given these many ways, radiation schemes can be approximated in various ways. For radiation schemes in medium to long range weather models, it makes sense to utilise more spectral bands to capture the complex shortwave radiative heating in the stratosphere while saving resources in the spatial and temporal dimensions by running the radiation scheme at coarser resolution and intermittently relative to the general model time stepping. Here we have chosen to work with the ACRANEB2 radiation scheme, which has been designed for short range weather models (Mašek et al., 2015; Geleyn et al., 2017). Detailed descriptions of the physics in ACRANEB2 have been made by Mašek et al. (2015) and Geleyn et al. (2017) for the shortwave and longwave radiation, respectively. More details about the dwarf can be found in Müller et al. (2017).

This dwarf has been entirely refactored which leads to massive speedups on CPU, GPU and KNL processors (Figure ??a). Optimisation of ACRANEB2 radiation dwarf from Poulsen and Berg (2017). (a) Time-to-solution relative to the baseline implementation on a SNB node for transt in the full acraneb2 dwarf. The baseline code performance on single nodes of different architectures is shown to the left. Bars in the middle show the single node performance of the refactored codes, and the right bars show the single core performance of the refactored codes. Using the Cray compiler on NVIDIA GPU K20x and

the PGI compiler on NVIDIA GPU P100, and the Intel compiler on Intel Broadwell (processor 2S E5-2699v4), KNL (type 7210) and Intel Sandy Bridge (2S E5-2680v1). Single-core performance is not sensible for the GPU. (b) Time-to-solution for the three different code bases on three different architectures. X is the Xeon target code. G is the GPU target code using the data structures as X, but split into seven chunks. GNM is the GPU target with transposed data structures as compared to X and G, reformulated power function and even more splits (into 12 chunks). To explore the full code optimisation potential three different code bases have been created which use different data layouts (Figure ??b). These results show that different data layouts can have a huge impact on performance. For more details about the work on optimising the radiation dwarf we refer to Poulsen and Berg (2017).

Further speedup beyond optimising the code can be achieved by computing some of the radiation at coarser resolution. Radiation schemes used in operational models consider only radiation in vertical columns. At resolutions around 1 km or finer we find that the accuracy of the radiation parameterisation degrades because it does not consider 3D effects like cloud shadows. Computing radiation at coarser resolution compared to the rest of the model has the potential to improve accuracy and gives us another significant speedup. Experiments with ACRANE2 indicate that we can expect a speedup of 10 times in realistic scenarios. A series of full 3D experiments with an operational model at DMI using 65 vertical levels and 800*600 grid points give a reduced execution time of about 25% for a cheap version of the radiation scheme and by about 70% if the radiation scheme is run in an expensive version every time step.

3.1 Cloud-microphysics scheme CLOUDSC

The CLOUDSC dwarf is the parametrization scheme for cloud and precipitation processes in the IFS, described by prognostic equations for cloud liquid water, cloud ice, rain, snow and a grid-box fractional cloud cover. The cloud scheme represents the sources and sinks of cloud and precipitation due to the major generation and destruction processes, including cloud formation by detrainment from cumulus convection, condensation, ice deposition, evaporation, hydrometeor collection, melting and freezing. The scheme is based on Tiedke (1993) but with an enhanced representation of the ice phase in clouds and precipitation. A multi-dimensional implicit solver is used for the numerical solution of the cloud and precipitation prognostic equations. A more detailed description of the formulation of the parametrization can be found in ECMWF (2015) with further discussion in Forbes and Tompkins (2011) and Forbes et al. (2011).

Independent of ESCAPE there has been some work on optimising CLOUDSC for GPUs (Xiao et al., 2017). This work added a hybrid MPI and OpenACC approach to the CLOUDSC dwarf and explored different optimisation methods for the GPU. Xiao et al. (2017) found that the performance is highly dependent on the size of the blocking as defined by the parameter NPROMA. The CPU version is most efficient with a relatively small size of the NPROMA block between 12 and 128 while the GPU version is most efficient with NPROMA over 10,000. In terms of pure computation time (Xiao et al., 2017) found the K80 GPU to be about two times faster than one CPU socket with 12 cores. If data movement between CPU and GPU is included the GPU was slightly slower than the CPU socket. We expect that having a faster interconnect like NVLink between CPU and GPU would improve the total GPU runtime significantly.

~~The work on CLOUDSC in ESCAPE focused on using GridTools we also worked on using domain specific languages via the CLAW DSL (Clement et al., 2018) for the cloud microphysics dwarf. In particular the use of the Single Column Abstraction (SCA), where physical parameterisations are defined solely in terms of a single horizontal column, enables domain scientists to define physics equations purely in terms of vertical dependencies without needing to account for parallelisation~~

5 issues. The CLAW DSL then inserts loops over the data-parallel horizontal dimension specific to the hardware architecture and programming model (OpenMP, OpenACC) via source-to-source translation, allowing multiple architectures to be targeted from a single source code. A GPU implementation of the CLOUDSC dwarf generated by automated source translation tools has been used to generate similar performance results to the ones presented by Xiao et al. (2017) on K80 GPUs via the OpenACC backend of the CLAW DSL.

10 ~~4 Comparison between different discretisation methods~~ Conclusions and outlook

~~Looking at a higher level comparison between different discretisation methods represented by our dwarfs, Kühnlein et al. (2018) compared the IFS-FVM (using MPDATA and GCR elliptic solver) with the spectral-transform formulation of IFS. Figure ?? shows close agreement between both model formulations for the baroclinic instability benchmark. Kühnlein et al. (2018) also demonstrate that IFS-FVM operates at a competitive computational performance compared to the operational spectral-transform IFS formulation at ECMWF.~~

15

~~Comparison of the surface pressure with (a) IFS-FVM and (b) the spectral-transform IFS for the baroclinic instability benchmark after 15 simulation days. The depiction illustrates comparable solutions with the different discretisation methods.~~

~~In terms of performance and energy efficiency of the entire forecast we compared ALARO (based on spectral transform using biFFT) with COSMO-EULAG (based on finite difference with MPDATA and elliptic solver). The results are shown in Figure ?. More details about this comparison can be found in Van Bever et al. (2018a). This comparison does not include many of the optimisations presented in this paper. As shown for the individual dwarfs both models are expected to still have significant potential for optimisations. We do not know if there will be a clear winner. Full model forecast run. Comparison of energy efficiency between ALARO (based on spectral transform using biFFT) and COSMO-EULAG (based on finite difference with MPDATA and elliptic solver). Log-log plot of the energy consumption vs. wall-clock time for the ALARO 2.5km and COSMO-EULAG 2.2km reference configuration. Only pure MPI jobs were simulated. The colours of the data points and added lines have the same meaning as in Figure 13.~~

20

25

~~Having energy vs. runtime plots like in Figure ?? is very important for NWP applications due to~~ The ESCAPE project has introduced the concept of Weather & Climate dwarfs as fundamental domain-specific building blocks into the weather and climate community. Their categorisation of computational and communication patterns has been extremely useful in further breaking down the complexity of weather prediction and climate models, and advancing their adaptation to future hardware. Prototype implementations of these dwarfs have been used to work on optimising them and using them for the purpose of benchmarking new computers. These included measures for verifying their scientific correctness, documentation and input from domain scientists. Our dwarfs are very well suited for optimisation and benchmarking, they are small enough to be fairly

30

~~easy to understand, and at the same time represent a significant part of the severe restrictions in terms of runtime and energy cost. As described before the forecast cannot start earlier because not all of the required observations would be available and it cannot finish later because it would not reach our customers on time. Finding the best compromise between energy consumption and runtime is therefore crucial.~~ performance of the whole weather prediction model.

5 5 Conclusions

~~This~~ The paper gives an overview of the ~~work done~~ dwarfs that were identified in the ESCAPE project. ~~The project aims at preparing weather prediction and climate models for new computing architectures towards exascale machines. The project introduced the concept of fundamental building blocks called dwarfs. Feedback from the European and international community at our dissemination workshops and at international conferences have shown that this work was well received and the entire~~ NWP community has a strong interest in pursuing the creation of dwarfs. Having dwarfs reduces the complexity of the code and enables HPC centres and hardware vendors to significantly improve, while illustrating with the spectral transform dwarf a detailed example of the optimisation cycle within ESCAPE. In ESCAPE-2 we further identify dwarfs in other Earth System components such as sea-ice and ocean models, that are crucial for the performance of ~~these smaller and self-contained portions of code.~~ Another important development is a common-coupled applications.

To avoid code duplication we used the data structure and mesh handling framework ~~as given~~ provided by Atlas. Atlas has been extended in ESCAPE to support limited area grids and DSL with GridTools.

The participating models and most prototype implementations are based on FORTRAN and all our optimisations can be incorporated into FORTRAN code including the use of CUDA functions on GPUs by calling C functions from FORTRAN. We have started to incorporate the obtained code optimisations into operations. ~~The first results of this effort look promising,~~ which gives us a significant speedup in the spectral transform used for postprocessing. Besides optimisations of the existing code, improved algorithms have been developed which are specifically targeted at improving performance on large scale systems. ~~We developed,~~ which include a multigrid preconditioner for the elliptic solver, ~~found~~ to reduce iteration counts in iterative solvers, a HEVI time-integration scheme with significantly improved stability, ~~explored~~ and alternative finite difference methods on the sphere ~~and explored~~ in the context of reducing global communications across large processor counts, and alternative solution procedures for spectral transforms at fixed energy cost, with FFTs and spherical harmonics realised on optical processors.

Code optimisations performed in the ESCAPE project targeted Intel CPUs, Intel Xeon Phi processors, NVIDIA GPUs and Optalysys optical processors. In NWP applications, the bottleneck in terms of performance is usually the memory bandwidth between processor and main memory. Having fast interconnect like NVLink and NVSwitch can provide a massive speedup. Having all of the processing units used by the simulation connected with such a fast interconnect is still a challenge. Using accelerators only for a small part of the code destroys a lot of the benefit in terms of overall cost if the CPUs are idle while the accelerators perform their computations. We either need to move a large part of the code to the accelerator (like in Fuhrer et al., 2018; Schalkwijk et al., 2015) or we need to overlap computations on the CPU with computations on the

accelerator. ~~Both of these options still require~~ (like in Dziekan et al., 2019). Applying this to the participating models still requires more work.

The best way to compare different processors that we could find for NWP applications is to compare the number of devices and the energy consumption required to reach a certain runtime. The spectral transform showed a huge benefit from using 5 ~~batched matrix multiplication which involves adding a large number of zero operations. This makes it impossible to compare any metrics that involve the number of operations performed by~~ Most of the work done in the ESCAPE project in terms of code optimisation focused on hardware specific optimisations. This makes the optimisations difficult to maintain on upcoming new hardware architectures. As an approach towards performance portability we implemented a DSL prototype for advection with MPDATA by using the GridTools framework, and a prototype of the processor.

10 ~~Our work on the radiation dwarf indicates that the changes between different architectures required to reach good performance can be very substantial.~~ cloud microphysics dwarf by using the CLAW DSL. DSLs are a very promising tool to enable good performance on multiple architectures while ~~still having one~~ maintaining a single code base. However, designing a domain specific language that is user friendly and at the same time close to hand tuned performance on each architecture is ~~still challenging and further work needs to be done.~~ challenging. This work will continue in the ESCAPE-2 project which started in 15 October 2018.

Comparing different methods requires to include all costs. Spectral transform ~~with semi-implicit, semi-Lagrangian time-integration~~ has always been considered as being poorly suited for large supercomputers due to the ~~large amount of communication per time-step and fairly large communicators~~ high volume of communication. Our work indicates that thanks to much larger time-steps ~~and better strong scaling than for halo communication~~ the overall communication cost is not necessarily worse than 20 for other methods. Again overlapping different parts of the model such that useful computation can be done while data is communicated is a way forward and needs to be high priority in future research.

When moving towards very high resolution global simulations of O(1km) or less and considering exascale computations on a variety of emerging HPC architectures, there is a continued interest and need in pursuing fundamentally different algorithmic approaches that simply do not communicate beyond a certain halo size while retaining all other favourable properties ~~of,~~ such as 25 ~~removing time-step size restrictions as in the semi-Lagrangian advection and with semi-implicit time-stepping (SISL) approach,~~ and such approaches are ~~to be~~ further investigated in ESCAPE2.

ESCAPE-2.

Code availability. The data structure framework Atlas is available at <https://github.com/ecmwf/atlas> under an Apache License 2.0. The Grid-Tools framework used as a domain specific language approach for MPDATA is available at <https://github.com/GridTools/gridtools> under a 30 BSD-3-Clause license. The CLAW DSL used for the cloud microphysics dwarf is available at <https://github.com/claw-project/claw-compiler> under a BSD-2-Clause license. Model codes developed at ECMWF are the intellectual property of ECMWF and its member states, and therefore the IFS code and the IFS-FVM code are not publicly available. Access to a reduced version of the IFS code may be obtained from ECMWF under an OpenIFS license (see <http://www.ecmwf.int/en/research/projects/openifs> for further information, last access: 28 May

2019). The ALARO and ALADIN codes, along with all their related intellectual property rights, are owned by the members of the ALADIN consortium and are shared with the members of the HIRLAM consortium in the frame of a cooperation agreement. This agreement allows each member of either consortium to license the shared ALADIN-HIRLAM codes to academic institutions of their home country for noncommercial research. Access to the codes of the ALADIN System can be obtained by contacting one of the member institutes or by submitting a request in the contact link below the page of the ALADIN website (<http://www.umr-cnrm.fr/aladin/>, last access: 28 May 2019) and the access will be subject to signing a standardised ALADIN-HIRLAM license agreement. The COSMO-EULAG model, along with all of its related intellectual property rights, is owned by the members of the COSMO consortium under the cooperation agreement. This agreement allows each member of the COSMO consortium to license the COSMO-EULAG code (<http://www.cosmo-model.org/content/consortium/licencing.htm>, last access: 28 May 2019) without fee to academic institutions of their home country for a noncommercial research. The code of the GRASS model is intellectual property of Météo-France and is not publicly available. The code of the ESCAPE dwarfs is intellectual property of ECMWF. A license for educational and non-commercial research can be obtained from ECMWF (see <http://www.hpc-escape.eu> for contact details, last access: 28 May 2019). For the GPU optimisation of the spectral transform dwarf we used the PGI compiler version 17.10, CUDA 9.0.176 and OpenMPI 2.1.3. For the GPU optimisation of the MPDATA dwarf we used the PGI compiler version 17.10, CUDA 9.2.88 and OpenMPI 2.1.3. CUDA includes the compiler nvcc used to compile the library function wrappers, the libraries themselves and the profiling tool nvprof which we used for profiling on the GPUs. For the CPU optimisation we used Intel compilers and libraries version 2018.1.163. This includes the compilers icc and ifort and the libraries mkl and mpi. The work on the optical processor Optalysys used the MATLAB software version 2017b.

5 Simple theoretical model for heterogenous computing

We derive in this appendix a simple theoretical model for cost of one time-step on a heterogenous supercomputer. For simplicity we assume that we have one type of accelerator with N_{acc} devices and our code consists of two parts: part A can only run on the CPUs whereas part B can run entirely on the accelerator. We start by deriving the number of CPU sockets N_{CPU} and accelerator devices N_{acc} if we have no overlap between computations on the CPU and the accelerator. In the second subsection we derive these numbers for the case of having perfect overlap between CPUs and accelerators. The section ends by using these numbers in a simple cost model and with a discussion of limitations of this simple cost model.

4.1 No overlap between CPU and accelerator

In this subsection we assume that we have no overlap between computations on the CPU and computations on the accelerator, i.e. for the total runtime of the model is given by

$$T = T_{\text{A,CPU}} + T_{\text{B,acc}},$$

where $T_{\text{A,CPU}}$ is the runtime of part A running on the CPU and $T_{\text{B,acc}}$ is the runtime of part B running on the accelerator. We can obtain these partial runtimes out of the serial runtimes by using Amdahl's law. For simplicity we assume that part A on the

CPU's and part B on the accelerator devices have the same proportion p that benefits from parallelisation:-

$$T_{A,CPU} = \left(1 - p + \frac{p}{N_{CPU}}\right) \tau_{A,CPU},$$

$$T_{B,acc} = \left(1 - p + \frac{p}{N_{acc}}\right) \tau_{B,acc},$$

where $\tau_{A,CPU}$ is the runtime of part A on one CPU socket and $\tau_{B,acc}$ is the runtime of part B on one accelerator device. The total

5 runtime T is a fixed given number in NWP applications. We can therefore solve T for N_{CPU} and obtain

$$N_{CPU} = \frac{N_{acc} p \tau_{A,CPU}}{N_{acc} (T + (p - 1) (\tau_{A,CPU} + \tau_{B,acc})) - \tau_{B,acc} p}.$$

For simplicity we assume that one CPU socket has the same hardware cost like one accelerator device. Under this assumption we can minimise the total cost by minimising the total number of devices-

$$N = N_{CPU} + N_{acc}.$$

10 Inserting (??) into (??) and minimising this expression as a function of N_{acc} gives us-

$$N_{acc} = \frac{p (\tau_{B,acc} + \sqrt{\tau_{A,CPU} \tau_{B,acc}})}{T + (p - 1) (\tau_{A,CPU} + \tau_{B,acc})}.$$

We now introduce the speed-up of the accelerator device by defining-

$$S_{acc} = \frac{\tau_{B,CPU}}{\tau_{B,acc}},$$

where $\tau_{B,CPU}$ is the runtime of the accelerator part of the code if a CPU version of this code is run on one CPU socket. We

15 denote the fraction of the code that is running on the accelerator with-

$$\alpha = \frac{\tau_{B,CPU}}{\tau},$$

where $\tau = \tau_{A,CPU} + \tau_{B,CPU}$ is the total runtime on one CPU socket. For $\alpha = 0$ the entire code would need to run on the CPU and for $\alpha = 1$ the entire code would run on the accelerator. This gives us-

$$\tau_{B,acc} = \alpha \tau / S_{acc},$$

$$20 \quad \tau_{A,CPU} = (1 - \alpha) \tau.$$

The variables τ , α and S_{acc} are given parameters for NWP applications.-

To make our results easier to interpret we substitute the parallelisation factor p with the strong scaling efficiency β when comparing N_{eff} devices with one single device:-

$$\beta = \frac{T_{p=1}}{T} = \frac{1}{N_{eff}(1 - p) + p}.$$

As a specific example we consider the case of Météo-France with the model AROME. Météo-France has two clusters with each containing 1800 bi-socket Broadwell nodes. AROME runs operationally on 180 nodes (360 Broadwell sockets). The runtime of the entire operational forecast is 30 minutes. The forecast range is 30 hours with 50s steps (2160 time-steps). This means that $T = 0.8s$. Through linear regression of real simulations with the AROME model we find that $\tau = 215.19s$ and the strong scaling efficiency would be $\beta = 53\%$ for $N_{\text{eff}} = 1000$. Using these values gives us the solid lines shown in Figure ?? if we plot N_{CPU} , N_{acc} and the sum of the two as a function of α . Number of CPU sockets or accelerator devices as a function of how much of the code runs on the accelerator. Solid lines show the results with no overlap between CPUs and accelerators as derived in Section ?? while dotted lines show the results with overlap as derived in Section ?. These results use the specific case of the AROME model at Météo-France including the experimentally through linear regression found strong scaling efficiency of $\beta = 53\%$ on $N_{\text{eff}} = 1000$ CPU sockets compared to one CPU socket. One accelerator device is assumed to be two times faster than one CPU socket for this figure, i.e. $S_{\text{acc}} = 2$.

4.1 Perfect overlap between CPU and accelerator

If we assume perfect overlap between computations on the CPUs and accelerators the number of devices required for our simulation are simply given by solving $T = T_{\text{A,CPU}}$ for N_{CPU} and solving $T = T_{\text{B,acc}}$ for N_{acc} . This gives us

$$N_{\text{CPU}} = \frac{p \tau_{\text{A,CPU}}}{T - \tau_{\text{A,CPU}} + p \tau_{\text{A,CPU}}},$$

$$N_{\text{acc}} = \frac{p \tau_{\text{B,acc}}}{T - \tau_{\text{B,acc}} + p \tau_{\text{B,acc}}}.$$

If we replace (??) with (??) and (??) with (??) we obtain the dotted curves in Figure ?. The difference between solid and dotted lines in Figure ? show the overhead caused by having idle devices in the case of no overlap.

4.1 Cost model

The energy consumption of the idle processors during one single time-step Δt is given by:

$$E_{\text{idle}} = T(N_{\text{CPU}} P_{0,\text{CPU}} + N_{\text{acc}} P_{0,\text{acc}}),$$

where $P_{0,*}$ is the power consumption of one device when it is idle. The additional energy consumption due to running the dwarfs would be during one single time-step Δt :

$$E_{\text{load}} = T_{\text{A,CPU}} N_{\text{CPU}} \Delta P_{\text{CPU}} + T_{\text{B,acc}} N_{\text{acc}} \Delta P_{\text{acc}},$$

where $\Delta P_* = P_* - P_{0,*}$ is the increase in energy consumption by not being idle. This gives us the total energy consumption per time-step

$$E_{\text{total}} = E_{\text{load}} + E_{\text{idle}}.$$

The total energy cost of the machine is therefore

$$C_{\text{total}} = \gamma_E E_{\text{total}},$$

where γ_E is the cost per energy.

Supercomputers at NWP centres are typically rented. If we denote the cost of renting one CPU socket with C_{CPU} and one accelerator device with C_{acc} we get for the cost of the hardware during one time-step

$$C_{\text{hw}} = T(N_{\text{CPU}}C_{\text{CPU}} + N_{\text{acc}}C_{\text{acc}}).$$

5 Therefore, we get for the total cost of one time-step

$$C_{\text{total}} = C_E + C_{\text{hw}}.$$

As a specific example we consider again the case of running the NWP model AROME at Météo-France. The monthly rent of the two clusters at Météo-France is 500k€. We measured that an idle node uses about 0.03kW, while a node running AROME uses 0.3kW. The cost of energy is approximately 0.15€/kWh. For simplicity we assume that one accelerator device
10 has the same energy consumption and renting cost like one CPU socket, i.e. $P_{0,\text{CPU}} = P_{0,\text{acc}} = 15\text{W}$, $\Delta P_{\text{CPU}} = \Delta P_{\text{acc}} = 135\text{W}$, $C_{\text{acc}} = C_{\text{CPU}} = 2.7 \times 10^{-5}/\text{s}$. Using these values in equation (??) gives us the results in Figure ??.

4.1 Limitations of the cost model

The reader needs to be aware of the simplifications used in deriving the simple cost model in this section. As described before the purpose of this cost model is to show that having no overlap between computations on the CPU and on the accelerator can
15 cause a significant overhead on hybrid machines. The cost model does not explicitly consider data transfer between CPU and accelerator. One could assume that the data transfer is included in the speedup of the accelerator but then the speedup would depend on α (i.e. how much of the model is running on the accelerator) and if all of the code ported to the accelerator is in one block resulting in one data transfer in every time step.

The cost model assumes that one CPU socket has the same energy consumption like one accelerator device and the code
20 achieves the same strong scaling efficiency on both. As stated before these assumptions might be wrong in reality. Same scaling efficiency will be difficult to achieve because accelerators usually require a lot of inherent parallelism inside each device.

Despite these simplifications, our cost model illustrates the overhead caused by not overlapping computations on CPU and accelerator. The reader needs to be aware that in order to make good use of accelerators we need to port a large part of the model to the accelerators and we need to invest in research on overlapping different computations.

25 *Author contributions.* Andreas Müller performed energy measurements for the spectral transform dwarf, implemented the optimised spectral transform for post-processing, worked on the implementation of the MPDATA dwarf, supported maintenance for all dwarf prototypes and assembled the paper. Willem Deconinck is the main developer of the data structure framework Atlas and contributed the work on improving and extending Atlas. He also was involved in creating, supporting and maintaining all of the dwarf prototypes. Christian Kühnlein and Piotr K. Smolarkiewicz are the main developers of IFS-FVM and contributed to the corresponding elliptic solver and MPDATA dwarfs. Gianmarco
30 Mengaldo was involved in the creation of the dwarfs and provided improvements to early drafts of the manuscript. Michael Lange and Valentin Clement provided DSL work for the cloud microphysics dwarf. Nils Wedi and Peter Bauer supervised the entire ESCAPE project as

coordinators. They were also involved in writing and improving the manuscript. Nils Wedi performed the measurements for the cost profiles in Figure 3. Michail Diamantakis implemented the semi-Lagrangian dwarf prototype. Sarah-Jane Lock supported the work on HEVI time integration methods. Mats Hamrud, Sami Saarinen and George Mozdzynski were involved in creating initial versions of the dwarf prototypes. Sami Saarinen also contributed the code for energy measurements on the Cray XC40 supercomputer. Daniel Thiemert provided project management for the ESCAPE project. Michael Ginton and Pierre Bénard implemented the GRASS model. Fabrice Voitus and Charles Colavolpe worked on HEVI time-stepping methods. Philippe Marguinaud and Yongjun Zheng performed simulations of communication costs. Joris Van Bever, Daan Degrauwe, Geert Smet and Piet Termonia provided energy and runtime measurements for ALARO, analysed the components of the RMI ensemble prediction system and implemented limited area support in Atlas. Kristian P. Nielsen, Bent H. Sass, Jacob W. Poulsen and Per Berg provided work on the radiation dwarf. Carlos Osuna and Oliver Fuhrer provided work on GridTools and were involved in adding GPU support to Atlas. Michael Baldauf contributed work on HEVI DG. Mike Gillard and Joanna Szmelter worked on improving the preconditioning for the elliptic solver. Enda O'Brien, Alastair McKinstry, Oisín Robinson, Parijat Shukla and Michael Lysaght provided the LAITRI dwarf and created initial CPU optimisations and OpenACC ports for all of the dwarf prototypes. Michał Kulczewski, Miłosz Ciznicki, Wojciech Piątek, Sebastian Ciesielski, Marek Błażewicz, Krzysztof Kurowski, Marcin Procyk, Paweł Sychala and Bartosz Bosak worked on performance modelling using the roofline model which allowed a better understanding of the experimental measurements. Zbigniew Piotrowski and Andrzej Wyszogrodzki compared different time-integration methods, supported optimisation of IFS-FVM and provided energy measurements of COSMO-EULAG. Erwan Raffin, Cyril Mazauric, David Guibert, Louis Douriez and Xavier Vigouroux provided the CPU optimisations for the spectral transform and MPDATA dwarf prototypes. Alan Gray and Peter Messmer contributed the GPU optimisations for the spectral transform and MPDATA dwarf prototypes. Alexander J. Macfaden and Nick New contributed the work on using optical processors for the spectral transform.

Acknowledgements. The ESCAPE project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No ~~671627-671627~~, see also <http://www.hpc-escape.eu/>. This work has been also supported by HPC resources provided by Poznan Supercomputing and Networking Center and corresponding research activities under the MAESTRO grant number DEC-2013/08/A/ST6/00296 in National Science Centre (NCN). An award of computer time was provided by the INCITE program. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE office of Science User Facility supported under contract DE-AC05-00OR22725.

References

- Asanović, K., Bodik, R., Catanzaro, B. C., Gebis, J. J., Husbands, P., Keutzer, K., Patterson, D. A., Plishker, W. L., Shalf, J., Williams, S. W., and Yelick, K. A.: The landscape of parallel computing research: a view from Berkeley, Tech. Rep. UCB/EECS-2006-183, EECS Department, University of California, Berkeley, <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-183.html>, 2006.
- 5 Asanović, K., Wawrzynek, J., Wessel, D., Yelick, K., Bodik, R., Demmel, J., Keaveny, T., Keutzer, K., Kubiawicz, J., Morgan, N., Patterson, D., and Sen, K.: A view of the parallel computing landscape, *Communications of the ACM*, 52, 56, <https://doi.org/10.1145/1562764.1562783>, 2009.
- Bauer, P., Thorpe, A., and Brunet, G.: The quiet revolution of numerical weather prediction, *Nature*, 525, 47–55, <https://doi.org/10.1038/nature14956>, 2015.
- 10 Bénard, P. and Glinon, M.: Circumventing the pole problem of reduced lat-lon grids with local schemes. Part I: analysis and model formulation, *Q.J.R. Meteorol. Soc.*, <https://doi.org/10.1002/qj.3509>, 2019.
- Clement, V., Ferrachat, S., Fuhrer, O., Lapillonne, X., Osuna, C. E., Pincus, R., Rood, J., and Sawyer, W.: The CLAW DSL, in: *Proceedings of the Platform for Advanced Scientific Computing Conference on - PASC '18*, ACM Press, <https://doi.org/10.1145/3218176.3218226>, 2018.
- 15 Colavolpe, C., Voitus, F., and Bénard, P.: RK-IMEX HEVI schemes for fully compressible atmospheric models with advection: analyses and numerical testing, *Q.J.R. Meteorol. Soc.*, 143, 1336–1350, <https://doi.org/10.1002/qj.3008>, 2017.
- Colella, P.: Defining software requirements for scientific computing, DARPA HPCS Presentation, 2004.
- Deconinck, W.: Development of the flexible data structure framework Atlas, Tech. rep., ESCAPE, <https://goo.gl/11RD6b>, 2017a.
- Deconinck, W.: Public release of Atlas library version under an open source license which is accelerator-enabled and has improved interoperability features, Tech. rep., ESCAPE, <https://goo.gl/dXaPoL>, 2017b.
- 20 Deconinck, W., Bauer, P., Diamantakis, M., Hamrud, M., Kühnlein, C., Maciel, P., Mengaldo, G., Quintino, T., Raoult, B., Smolarkiewicz, P. K., and Wedi, N. P.: Atlas - A library for numerical weather prediction and climate modelling, *Comp. Phys. Communications*, 220, 188–204, <https://doi.org/10.1016/j.cpc.2017.07.006>, 2017.
- Douriez, L., Gray, A., Guibert, D., Messmer, P., and Raffin, E.: Performance report and optimized implementations of Weather & Climate dwarfs on multi-node systems, Tech. rep., ESCAPE, <http://tiny.cc/fzx36y>, 2018.
- 25 Dziekan, P., Waruszewski, M., and Pawlowska, H.: University of Warsaw Lagrangian Cloud Model (UWLCM) 1.0: a modern Large-Eddy Simulation tool for warm cloud modeling with Lagrangian microphysics, *Geoscientific Model Development Discussions*, pp. 1–26, <https://doi.org/10.5194/gmd-2018-281>, 2019.
- ECMWF, R.-D.: IFS documentation, <https://www.ecmwf.int/en/forecasts/documentation-and-support/changes-ecmwf-model/ifs-documentation>, 2015.
- 30 Eisenstat, S. C., Elman, H. C., and Schultz, M. H.: Variational iterative methods for nonsymmetric systems of linear equations, *SIAM Journal on Numerical Analysis*, 20, 345–357, <https://doi.org/10.1137/0720023>, 1983.
- Feng, W.-C., Lin, H., Scogland, T., and Zhang, J.: OpenCL and the 13 dwarfs, in: *Proceedings of the third joint WOSP/SIPEW international conference on Performance Engineering - ICPE '12*, ACM Press, <https://doi.org/10.1145/2188286.2188341>, 2012.
- 35 Flamm, K.: Measuring Moore’s Law: Evidence from Price, Cost, and Quality Indexes, Tech. rep., <https://doi.org/10.3386/w24553>, 2018.
- Forbes, R. and Tompkins, A.: An improved representation of cloud and precipitation, *ECMWF Newsletter No. 129*, pp. 13–18, 2011.
- Forbes, R., Tompkins, A., and Untch, A.: A new prognostic bulk microphysics scheme for the IFS, *ECMWF Tech. Memo. No. 649*, 2011.

- Fuhrer, O., Chadha, T., Hoefler, T., Kwasniewski, G., Lapillonne, X., Leutwyler, D., Lüthi, D., Osuna, C., Schär, C., Schulthess, T. C., and Vogt, H.: Near-global climate simulation at 1km resolution: establishing a performance baseline on 4888 GPUs with COSMO 5.0, *Geosci. Model Dev.*, 11, 1665–1681, <https://doi.org/10.5194/gmd-11-1665-2018>, 2018.
- Geleyn, J.-F., Mašek, J., Brožková, R., Kuma, P., Degrauwe, D., Hello, G., and Pristov, N.: Single interval longwave radiation scheme based on the net exchanged rate decomposition with bracketing, *Q.J.R. Meteorol. Soc.*, 143, 1313–1335, <https://doi.org/10.1002/qj.3006>, 2017.
- 5 Grinton, M. R. and Bénard, P.: Circumventing the pole problem of reduced lat-lon grids with local schemes. Part II: validation experiments, *Q.J.R. Meteorol. Soc.*, <https://doi.org/10.1002/qj.3495>, 2019.
- Hamrud, M.: Report from IFS scalability project, Technical Memorandum 616, ECMWF, 2010.
- Hortal, M.: The development and testing of a new two-time-level semi-Lagrangian scheme (SETTLS) in the ECMWF forecast model, *Tech. Rep. 292*, Eur. Cent. For Medium-Range Weather Forecasts, Reading, UK, <https://doi.org/10.1256/00359000260247417>, 1999.
- 10 Johnston, B. and Milthorpe, J.: Dwarfs on Accelerators, in: Proceedings of the 47th International Conference on Parallel Processing Companion - ICPP '18, ACM Press, <https://doi.org/10.1145/3229710.3229729>, 2018.
- Kaltofen, E. L.: The seven dwarfs of symbolic computation, in: *Texts & Monographs in Symbolic Computation*, pp. 95–104, Springer Vienna, https://doi.org/10.1007/978-3-7091-0794-2_5, 2011.
- 15 Katzav, J. and Parker, W. S.: The future of climate modeling, *Climatic Change*, 132, 475–487, <https://doi.org/10.1007/s10584-015-1435-x>, 2015.
- Krommydas, K., chun Feng, W., Antonopoulos, C. D., and Bellas, N.: OpenDwarfs: Characterization of Dwarf-Based Benchmarks on Fixed and Reconfigurable Architectures, *Journal of Signal Processing Systems*, 85, 373–392, <https://doi.org/10.1007/s11265-015-1051-z>, 2015.
- Kühnlein, C. and Smolarkiewicz, P. K.: An unstructured-mesh finite-volume MPDATA for compressible atmospheric dynamics, *J. Comput. Phys.*, 334, 16–30, <https://doi.org/10.1016/j.jcp.2016.12.054>, 2017.
- 20 Kühnlein, C., Deconinck, W., Klein, R., Malardel, S., Piotrowski, Z. P., Smolarkiewicz, P. K., Szmelter, J., and Wedi, N. P.: FVM 1.0: A nonhydrostatic finite-volume dynamical core formulation for IFS, *Geosci. Model Dev. Discuss.*, <https://doi.org/10.5194/gmd-2018-237>, 2018.
- Kühnlein, C., Deconinck, W., Klein, R., Malardel, S., Piotrowski, Z. P., Smolarkiewicz, P. K., Szmelter, J., and Wedi, N. P.: FVM 1.0: a nonhydrostatic finite-volume dynamical core for the IFS, *Geoscientific Model Development*, 12, 651–676, <https://doi.org/10.5194/gmd-12-651-2019>, 2019.
- 25 Lawrence, B. N., Rezny, M., Budich, R., Bauer, P., Behrens, J., Carter, M., Deconinck, W., Ford, R., Maynard, C., Mullerworth, S., Osuna, C., Porter, A., Serradell, K., Valcke, S., Wedi, N., and Wilson, S.: Crossing the chasm: how to develop weather and climate models for next generation computers?, *Geosci. Model Dev.*, 11, 1799–1821, <https://doi.org/10.5194/gmd-11-1799-2018>, 2018.
- 30 Lock, S.-J., Wood, N., and Weller, H.: Numerical analyses of Runge-Kutta implicit-explicit schemes for horizontally explicit, vertically implicit solutions of atmospheric models, *Q.J.R. Meteorol. Soc.*, 140, 1654–1669, <https://doi.org/10.1002/qj.2246>, 2014.
- Macfaden, A. J., Gordon, G. S. D., and Wilkinson, T. D.: An optical Fourier transform coprocessor with direct phase determination, *Scientific Reports*, 7, <https://doi.org/10.1038/s41598-017-13733-1>, 2017.
- Magnusson, L., Bidlot, J.-R., Bonavita, M., Brown, A., Browne, P., Chiara, G. D., Dahoui, M., Lang, S. T. K., McNally, T., Mogensen, K. S., Pappenberger, F., Prates, F., Rabier, F., Richardson, D. S., Vitart, F., and Malardel, S.: ECMWF activities for improved hurricane forecasts, *Bulletin of the American Meteorological Society*, <https://doi.org/10.1175/bams-d-18-0044.1>, 2018.
- 35 Mašek, J., Geleyn, J.-F., Brožková, R., Giot, O., Achom, H. O., and Kuma, P.: Single interval shortwave radiation scheme with parameterized optical saturation and spectral overlaps, *Q.J.R. Meteorol. Soc.*, 142, 304–326, <https://doi.org/10.1002/qj.2653>, 2015.

- Mazauric, C., Raffin, E., and Guibert, D.: Report on recommendations and specifications for data scope analysis tools, Tech. rep., ESCAPE, <https://goo.gl/61DZxq>, 2017a.
- Mazauric, C., Raffin, E., Vigouroux, X., Guibert, D., Macfaden, A., Poulsen, J., Berg, P., Gray, A., and Messmer, P.: Performance report and optimized implementation of weather & climate dwarfs on GPU, MIC and Optalysys optical processor, Tech. rep., ESCAPE, <https://goo.gl/J8zd1i>, 2017b.
- 5
- Mengaldo, G.: Batch 1: definition of several weather & climate dwarfs based on established algorithms and motifs relevant to NWP, provision of prototype implementations and dissemination to other WPs, Tech. rep., ESCAPE, <http://goo.gl/s65ojl>, 2016.
- Mengaldo, G., Wyszogrodski, A., Diamantakis, M., Lock, S.-J., Giraldo, F., and Wedi, N. P.: Current and emerging time-integration strategies in global numerical weather and climate prediction, *Arch. Computat. Methods Eng.*, pp. 1–22, <https://doi.org/10.1007/s11831-018-9261-8>, 2018.
- 10
- Messer, O. E. B., D’Azevedo, E., Hill, J., Joubert, W., Berrill, M., and Zimmer, C.: MiniApps derived from production HPC applications using multiple programming models, *Int. J. High Performance Computing Applications*, 32, 582–593, <https://doi.org/10.1177/1094342016668241>, 2016.
- Michalakes, J., Govett, M., Benson, R., Black, T., Juang, H., Reinecke, A., and Skamarock, B.: AVEC Report: NGGPS Level-1 Benchmarks and software evaluation, Tech. Rep. TN-484, NCAR, Boulder, US, <https://www.weather.gov/media/sti/nggps/AVECLevel1BenchmarkingReport0820150602.pdf>, 2015.
- 15
- Mozdzyński, G., Hamrud, M., and Wedi, N. P.: A Partitioned Global Address Space implementation of the European Centre for Medium Range Weather Forecasts Integrated Forecasting System, *Int. J. High Perform. Comput. Appl.*, 29, 261–273, <https://doi.org/10.1177/1094342015576773>, 2015.
- 20
- Müller, A., Gillard, M., Nielsen, K. P., and Piotrowski, Z.: Batch 2: definition of novel weather & climate dwarfs, provision of prototype implementations and dissemination to other WPs, Tech. rep., ESCAPE, <https://goo.gl/zoCtck>, 2017.
- Müller, A., Kopera, M. A., Marras, S., Wilcox, L. C., Isaac, T., and Giraldo, F. X.: Strong scaling for numerical weather prediction at petascale with the atmospheric model NUMA, *Int. J. High Perform. Comput. Appl.*, <https://doi.org/10.1177/1094342018763966>, 2018.
- Müller, A., Kopera, M. A., Marras, S., Wilcox, L. C., Isaac, T., and Giraldo, F. X.: Strong scaling for numerical weather prediction at petascale with the atmospheric model NUMA, *Int. J. High Perform. Comput. Appl.*, <https://doi.org/10.1177/1094342018763966>, 2018.
- 25
- Neumann, P., Düben, P., Adamidis, P., Bauer, P., Brück, M., Kornbluh, L., Klocke, D., Stevens, B., Wedi, N., and Biercamp, J.: Assessing the scales in numerical weather and climate predictions: will exascale be the rescue?, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 377, 20180148, <https://doi.org/10.1098/rsta.2018.0148>, 2019.
- Osuna, C.: Report on the performance portability demonstrated for the relevant weather & climate dwarfs, Tech. rep., ESCAPE, <https://goo.gl/tbEiAu>, 2018.
- Palmer, T.: Climate forecasting: Build high-resolution global climate models, *Nature*, 515, 338–339, <https://doi.org/10.1038/515338a>, 2014.
- 30
- Phillips, S. C., Engen, V., and Papay, J.: Snow White Clouds and the Seven Dwarfs, in: 2011 IEEE Third International Conference on Cloud Computing Technology and Science, IEEE, <https://doi.org/10.1109/cloudcom.2011.114>, 2011.
- Poulsen, J. W. and Berg, P.: Tuning the implementation of the radiation scheme ACRANE2, Tech. rep., DMI report 17-22, http://www.dmi.dk/fileadmin/user_upload/Rapporter/TR/2017/SR17-22.pdf, 2017.
- Robinson, O., McKinstry, A., and Lysaght, M.: Optimization of IFS subroutine LAITRI on Intel Knights Landing, Tech. rep., Prace White Papers - Evaluations on Intel MIC, 2016.
- 35
- Schalkwijk, J., Jonker, H. J. J., Siebesma, A. P., and Meijgaard, E. V.: Weather Forecasting Using GPU-Based Large-Eddy Simulations, *Bulletin of the American Meteorological Society*, 96, 715–723, <https://doi.org/10.1175/bams-d-14-00114.1>, 2015.

- Schulthess, T. C., Bauer, P., Wedi, N., Fuhrer, O., Hoefler, T., and Schär, C.: Reflecting on the Goal and Baseline for Exascale Computing: A Roadmap Based on Weather and Climate Simulations, *Computing in Science & Engineering*, 21, 30–41, <https://doi.org/10.1109/mcse.2018.2888788>, 2019.
- 5 Scott, R. K., Harris, L. M., and Polvani, L. M.: A test case for the inviscid shallow-water equations on the sphere, *Q.J.R. Meteorol. Soc.*, 142, 488–495, <https://doi.org/10.1002/qj.2667>, 2015.
- Shukla, J., Palmer, T. N., Hagedorn, R., Hoskins, B., Kinter, J., Marotzke, J., Miller, M., and Slingo, J.: Toward a New Generation of World Climate Research and Computing Facilities, *Bulletin of the American Meteorological Society*, 91, 1407–1412, <https://doi.org/10.1175/2010bams2900.1>, 2010.
- Smolarkiewicz, P., Deconinck, W., Hamrud, M., Kühnlein, C., Mozdzyński, G., Szmelter, J., and Wedi, N. P.: A finite-volume module for simulating global all-scale atmospheric flows, *J. Comput. Phys.*, 314, 287–304, <https://doi.org/10.1016/j.jcp.2016.03.015>, 2016.
- 10 Smolarkiewicz, P. K. and Margolin, L. G.: MPDATA: A finite difference solver for geophysical flows, *J. Comput. Phys.*, 140, 459–480, <https://doi.org/10.1006/jcph.1998.5901>, 1998.
- Smolarkiewicz, P. K. and Szmelter, J.: MPDATA: an edge-based unstructured-grid formulation, *J. Comput. Phys.*, 206, 624–649, <https://doi.org/10.1016/j.jcp.2004.12.021>, 2005.
- 15 Szmelter, J. and Smolarkiewicz, P. K.: An edge-based unstructured mesh discretisation in geospherical framework, *J. Comput. Phys.*, 229, 4980–4995, <https://doi.org/10.1016/j.jcp.2010.03.017>, 2010.
- Temperton, C., Hortal, M., and Simmons, A.: A two-time-level semi-Lagrangian global spectral model, *Q.J.R. Meteorol. Soc.*, 127, 111–127, <https://doi.org/10.1002/qj.49712757107>, 2001.
- Thomas, S. J., Hacker, J. P., Smolarkiewicz, P. K., and Stull, R. B.: Spectral preconditioners for nonhydrostatic atmospheric models, *Mon. Weather Rev.*, 131, 2464–2478, [https://doi.org/10.1175/1520-0493\(2003\)131<2464:spfnam>2.0.co;2](https://doi.org/10.1175/1520-0493(2003)131<2464:spfnam>2.0.co;2), 2003.
- 20 Tiedke, M.: Representation of clouds in large-scale models, *Mon. Weather Rev.*, 121, 3040–3061, [https://doi.org/10.1175/1520-0493\(1993\)121<3040:rocils>2.0.co;2](https://doi.org/10.1175/1520-0493(1993)121<3040:rocils>2.0.co;2), 1993.
- Van Bever, J., McFaden, A., Piotrowski, Z., and Degrauwe, D.: Report on energy-efficiency evaluation of several NWP model configurations, Tech. rep., ESCAPE, <https://goo.gl/E8mHha>, 2018a.
- 25 Van Bever, J., Smet, G., and Degrauwe, D.: Report on workflow analysis for specific LAM applications, Tech. rep., ESCAPE, <https://goo.gl/oRJaHx>, 2018b.
- Wallemacq, P., UNISDR, and CRED: Economic Losses, Poverty and Disasters 1998-2017, <https://doi.org/10.13140/rg.2.2.35610.08643>, 2018.
- Wedi, N. P., Hamrud, M., and Mozdzyński, G.: A Fast Spherical Harmonics Transform for Global NWP and Climate Models, *Mon. Weather Rev.*, 141, 3450–3461, <https://doi.org/10.1175/mwr-d-13-00016.1>, 2013.
- 30 Wedi, N. P., Bauer, P., Deconinck, W., Diamantakis, M., Hamrud, M., Kühnlein, C., Malardel, S., Mogensen, K., Mozdzyński, G., and Smolarkiewicz, P.: The modelling infrastructure of the Integrated Forecasting System: recent advances and future challenges, Tech. Rep. 760, Eur. Cent. For Medium-Range Weather Forecasts, Reading, UK, 2015.
- Wehner, M. F., Oliker, L., Shalf, J., Donofrio, D., Drummond, L. A., Heikes, R., Kamil, S., Kono, C., Miller, N., Miura, H., Mohiyuddin, M., Randall, D., and Yang, W.-S.: Hardware/software co-design of global cloud system resolving models, *Journal of Advances in Modeling Earth Systems*, 3, <https://doi.org/10.1029/2011ms000073>, 2011.
- World Economic Forum: The 2019 Global Risks Report, <https://www.weforum.org/reports/the-global-risks-report-2019>, 2019.

- Wu, J., Wyckoff, P., and Panda, D.: High performance implementation of MPI derived datatype communication over InfiniBand, in: 18th International Parallel and Distributed Processing Symposium, 2004. Proceedings., IEEE, <https://doi.org/10.1109/ipdps.2004.1302917>, 2004.
- Wyrzykowski, R., Szustak, L., and Rojek, K.: Parallelization of 2D MPDATA EULAG algorithm on hybrid architectures with GPU accelerators, *Parallel Computing*, 40, 425–447, <https://doi.org/10.1016/j.parco.2014.04.009>, 2014a.
- Wyrzykowski, R., Szustak, L., Rojek, K., and Tomas, A.: Towards efficient decomposition and parallelization of MPDATA on hybrid CPU-GPU cluster, in: *Large-Scale Scientific Computing*, pp. 457–464, Springer Berlin Heidelberg, https://doi.org/10.1007/978-3-662-43880-0_52, 2014b.
- Xiao, H., Diamantakis, M., and Saarinen, S.: An OpenACC GPU adaptation of the IFS cloud microphysics scheme, ECMWF Tech. Memo. No. 805, 2017.
- Zheng, Y. and Marguinaud, P.: Simulation of the Performance and Scalability of MPI Communications of Atmospheric Models running on Exascale Supercomputers, *Geosci. Model Dev.*, pp. 1–34, <https://doi.org/10.5194/gmd-2017-301>, 2018.